The Logic Of Ignorance: Towards and Inverted Epistemology

By Nathan Anecone

Preliminaries 13

Terminology and Taxonomy 13

General Framework 20

The Logical Form of Ignorance 35

Introduction 35

Positive versus Negative Ignorance 40

Complex versus Simple Ignorance 46

Res Ignorans: the Substantiality of Ignorance 50

The Principle of Continuity of Ignorance 59

Structural Ignorance: Ignorance as Cause 70

Epistemological Incompleteness 76

The Epistemic Universe 80

The Unknowable as Thing in Itself 82

Ignorance and Information Processing 87

Introduction 87

The Importance of Uncertainty 88

Blurry Models and the Perception of Uncertainty 93

Thinking as Uncertainty Reduction 103

Decision-Making under Pervasive Ignorance 106

Imperfect-Information Games and Belief Formation 113

Belief Revision as Bayesian Nets 115

Laplace's Demon: The Cosmic Know-It-All

118
Review so far 123
The Interaction between knowledge and Ignorance
127
Introduction 127
The Unspecified Knower 134
Recursive and Reciprocal Relationships 143
The Fourfold Ordered Pairs of Ignorance-Knowledge Relations
146
The Paradox of Unconscious knowledge 149
knowledge Graphs and Ignorance Graphs
155
Adversarial Relationships 163
Uncertainty-Proofing 167
Eliminative Cognition as Structural Ignorance 171
Trial and Error as Primitive Uncertainty Reduction
171
Attention as Information Exclusion 176
Subtractive knowledge 179
Unsupervised Learning as Cognitive Bootstrapping
181
Part II, Section I: Physical and Computational Constraints on Knowability
185
Introduction 185
The Physics of the Unknown 190
The Informational Turn 190
A Primer on the Second Law of Thermodynamics
191
The Fundamental Role of Entropy in Cognition
193
The Principle of Maximum Entropy as the Physical Representation of Ignorance

The Principle of Maximum Entropy as the Physical Representation of Ignorance
The Principle of Maximum Entropy as the Physical Representation of Ignorance 198
The Supervenience of Cognition on Entropy 204
The Temporality of Ignorance 207
Hot Thought: Maxwell's Demon, the Szilard engine, and Landauer's Principle 212
knowledge Gradients 218
Quantum Mechanical Analogies for Belief Formation 222
Possibility, Uncertainty and the Path Integral Formulation
Possibility, Uncertainty and the Path Integral Formulation
Possibility, Uncertainty and the Path Integral Formulation 224
Physical Limits on Information Access: Black Holes, Light Cones and the Speed of Light 231
Hidden Variable Theories 240
Part II Sectiion II: Logical and Computational Limits 243
Introduction 243
The Limits of Computability and Cognitive Closure 247
Cognitive Closure or Universal Computation? 247
a: Incompleteness Theorems and Self-Referential Undecidability
a: Incompleteness Theorems and Self-Referential Undecidability
a: Incompleteness Theorems and Self-Referential Undecidability 252
b: Formal Definition: 263
Part III: The Modal Logic of Ignorance 267

Introduction 267
The Many Worlds of the Unknown 268
Introduction 268
The Conceptual Order of Possibility 269
Modal Neg-Epistemic Relations for Knowability 277
Alethia and the Order of Discovery Space 284
The Scaling of Ignorance to Complexity 291
The Problem of Induction from an Ignorance First Perspective 299
Introduction 299
The Conditioning of Induction by Structural Ignorance 300
Part IV : Ignorance In A State of Nature 316
Introduction 316
Agnoiology Naturalized 317
The Brain as the Empirical Object of Ignorance 317
The Open-Ended Brain 320
Senesence as a Manifestation of Structural Ignorance 328
Towards A Theory of Cognitive Bias 332
Evolution as an Epistemic Concern 340
Immunology: The Defense Against The Microbial Unknown 344
Does the Universe Learn? 350
Part V : Homo insapiens: The All-Too-Human Implications and Applications of the Theory of Ignorance 355
Introduction 355
Ignorance Generation 356
Agniogenesis: the Production of Ignorance

356
The Logic of Lying 365
On the Occasionally Destructive Power of Truth
370
Information Overload: The Overwhelmingness of the Unknown
375
The Psychology of Ignorance 379
Introduction 379
Consciousness as the Expression of Information Entropy
380
Emotionality 384
A Note on Faith 387
The Nature of Opinion 390
knowledge versus Opinion: A Matter of Degree or Kind?
knowledge versus Opinion: A Matter of Degree or Kind?
390
Willful Ignorance as The Will to Believe
400
Cultural, Artifactual and Generational Ignorance
405
The notion of ignorance production lends itself naturally to the idea that ignorance... 405
Part VII: The Theory of Inquiry 415
Introduction 415
The Inquisitive Stance: Ignorance as Intentional Object
416
Hintikka's Socratic Epistemology 442
The Ignorance-first Nature of Inquiry 460
The Inquisitive Stance: Ignorance as Intentional Object
460
Hintikka's Socratic Epistemology 487
Mechanisms of Investigation 505

Testing, Exploration and Probing 505
The Scientific Method as an Instrument of Ignorance Management 519
The Philosophy of Courts 523
The End or Beginning of Science? 527
Conclusion: Ignoramus et Ignorabimus? 530

Note:

I began writing The Logic of Ignorance in 2019. I would have never gotten as far as I did if not for a stroke of luck in the form of the Covid lockdowns and a spate of unemployment. In the ensuing years, life kept getting in the way, and I kept holding out for a time when I could properly complete it. I hired a philosophy PHD off of UpWork who convinced me that large tracts of it needed to be revised and rethought, and the thought of doing all all the editorial work that had to be done myself bored me half to death. (We now live in an age of AI, which could help with some of that work, but even then I don't think I'd fully escape it.)

For years I held out on the belief that I would eventually find the time and the resources to properly complete this book. I've since decided that time will never come, and if it did, I'd probably move onto other things. Much of what I've written here has been lost to that pervasive form of ignorance we call forgetfulness.

So I've decided to release it onto the internet, warts and all. You will notice that the citations are a mess, most of them are in the form of "inline" links posted in place, and some might be referred to only in passing. I never got to putting together a proper bibliography.

I stand by this work and most of its claims, although some are quite outlandish and poorly articulated. The amount of ignorance in the world is as constant as ever, so I like to think the book contains a perennial relevance.

-Nathan Anecone, 2025

Foreword

The tradition of epistemology is incomplete. That's not to say it's necessarily wrong. It only addresses at most a third of the complete understanding. That's because it has failed in its conception of the role of ignorance in conditioning what is known and knowable.

Ignorance is fundamental. Arguably it is central to our lives, more central than even knowledge because much knowledge is essentially a means to deal with the ignorance remaining. And there is quite a lot of it remaining, much more, as I will attempt to precisely demonstrate, than we will ever have of knowledge.

Ignorance cannot be written off as a mere negation or absence of knowledge. Merely supposing it to be the lack of knowledge alone, a void or emptiness, firstly obscures the major point about the *interestingness* of ignorance in our lives and its conceptual richness. Secondly it's an expression of ignorance itself to believe ignorance to be a void where knowledge could be. Ignorance weaves into our beliefs about the world, imitating the character of knowledge yet representing a falsified picture of affairs such as a delusion or hallucination, to name just two instances of what I call positive ignorance. This is ignorance that is in some way *built up* and representational.

In essence, ignorance has a describable logical structure which is diametrically at odds with the assumption that it is not knowledge or an absence thereof which would entail that it has no form at all. These theoretical underpinnings of the structure of ignorance are the primary concern, and the quest to understand them leads to startling conclusions.

Epistemology traditionally conceived has doubled up on its misunderstanding of ignorance as the negation of positive knowledge, its ignorance of ignorance. Not only does it misunderstand or ignore what ignorance is, it is ignorant of being wrong about it and that ignorance is fundamental to an adequate understanding of knowledge. Much progress has to be made to understand the complexities of ignorance and how it interacts with positive knowledge.

If we are to truly understand how we know it's not enough to stop at knowledge. There is more to knowledge than the traditional concerns of epistemology. For completenesses' sake, knowledge must be balanced by ignorance and certainty uncertainty.

Like knowledge and certainty, ignorance and uncertainty comprise the two faces of a Janus coin. Their relationship is integral to one another. Therefore the object here is as much uncertainty as ignorance.

To satisfy the requirements of this project we must also identify a means to relate the ignorance and knowledge. This is to be found in the theory of inquiry. Mechanisms of discovery, investigation and exploration directly interfaces knowledge and ignorance. Though not stressed with as much emphasis as ignorance in this book, the theory of inquiry is equally relevant to the task of a complete epistemology.

To these ends we will draw from an array of sources and concepts, from philosophy, neuroscience, artificial intelligence and computing, information theory, everyday life, mathematics, logic, and physics. Together, the synthesis of the theories of knowledge, inquiry and ignorance yields many insights into a wide variety of problems.

The goal of this work is ambitious. Like any ambitious goal it is more likely to fall short of its aim rather than meet or exceed it. Even if that ends up being the case, I see it more as vindication of my premises rather than failure to proof them. If I fail to draw obvious connections or some of my arguments are plainly incorrect, I am merely proving my own point in an indirect sort of way. Some might call that cheating. In any event it would be cute if a book about ignorance, that sees ignorance nearly everywhere, should be exempt from the very property it ascribes to everything else.

In the course of writing this, I came across the book *Understanding Ignorance* by Daniel DeNicola. I was immediately struck by extensive parallels. The convergence between this work and that of DeNicola's, to my mind, is no mere coincidence. It points to the underlying substantial reality of ignorance. Unlike DeNicola's book, it pays attention less to the deeply human aspects of ignorance such as its ethical implications, although these human concerns are

addressed. Rather, it seeks to present a theoretical framework for ignorance in its most abstracted logical form. For the reader looking for a more accessible and humane portrayal of ignorance that taps into the same line of thinking as what's presented here, the author recommends professor DeNicola's book.

The present work assumes a fairly high level of background knowledge from the reader, although some chapters and sections are less technical than others.

In 2002 Timothy Williamson published an influential book *knowledge and its Limits* in which he proposed the concept of knowledge-first epistemology. In contradistinction to conventional wisdom, the knowledge-first account holds that knowledge is an unanalyzable, primitive concept, irreducible to constitutive atomic concepts such as belief, truthfulness, or justification as is conventionally assumed. knowledge stands on its own, and is understandable only from its own conditions and level of description (Williamson, 2002). While I have some gripes with this view, and prefer the reliable if contested old school view that knowledge has something to do with belief, truth, and justification, this is not the point. In a similar vein, the aim here is to propose an inverted, ignorance-first epistemology, a view which conceives of ignorance as, in loose and slightly misleading terms, an (un)equal and opposite (anti)epistemic state that counterbalances epistemic states. Ignorance truly is primitive, whole and uniform, and cannot be reduced to more elementary concepts. Metaphorically speaking, like zero, if you add, subtract, or multiply ignorance to itself you still get zero (ignorance). If you apply arithmetic to knowledge, however, you change its value, you get more or less of it. Moreover, in a slight twist I wish to emphasize the double meaning of "first", in both Williamson's sense of conceptual independence from other mental concepts and in the temporal sense that ignorance comes at the beginning of a sequence of events that terminates in knowledge. Ignorance unambiguously comes before knowledge. We are always ignorant of a fact before we come to know it, born ignorant except for genetic predispositions, etc. So if ignorance is not something real that relates to knowledge, it begs the question of how knowledge comes about at all if it is not somehow self-creating and self-causing and therefore pre-existent. knowledge does not pop out of nowhere. Knowledge in some meaningful way grows out of

ignorance by emerging from it through the process of inquiry that draws out discovered facts which were at one point occluded by ignorance, collapsing uncertainty and possibilities into a solidified engram or cognitive structure. Therefore it is impossible to truly understand knowledge realistically unless we understand how it operates in temporal and logical relation to ignorance, which, through a kind of oppositional tension, polarizes and gives direction to knowledge.

One does not need to subscribe to the knowledge-first view to map some key parallels from it to an ignorance-first view. Indeed, it is inaccurate to suppose a perfect symmetry between knowledge and ignorance because ignorance is a broader category over and above the mere negation of knowledge. It is not the logical complement of knowledge plain and simple, although there is a subset of ignorance at its simplest which is. The knowledge-first account could be false but the ignorance-first account still hold true independent of the fact. The goal is not to impersonate the knowledge-first story in the costume of ignorance. If anything ignorance-first strikes down the knowledge-first view, because it suggests that knowledge can be analyzed into more primitives, namely information states. In the process of becoming knowledge, knowledge is as it were built piecemeal out of information extracted from the gradual unveiling of the unknown by inquiry as we are drawn to eliminate our awareness of what we believe we don't know. In the latter parts of the book this point will be made abundantly clear: inquiry does not elicit knowledge from a source but only information, because only information can be said to exist external to the subject. knowledge does not enter the brain though photons hitting the retina or sound waves vibrating the ear drum; at some point these raw signals translate into knowledge through some act of cognitive self-construction, but at what point does information transition into the knowledge? Because it is extremely difficult to clearly distinguish the transduction of physical stimuli into neural activity into the resulting knowledge, one must assume that knowledge is somehow made out of information. This idea will be elaborated in due course.

Because ignorance is the inescapable backdrop of knowledge and thinking, it is curiously useful to view it as a kind of logical "negative substance" like dark

matter or empty space as opposed to regular "masses" of knowledge. Both must be specified to define the epistemic universe. This negative substance lacks observable, inspectable properties in the traditional sense. It is elusive, lacking content yet holding form; it can somehow be categorized into a typology but lacks all concrete properties. The confusion here however is that ignorance, while lacking in tangible attributes, has a logical structure that can be traced out through careful detective work. Its properties are not manifest but are logical consequences of reasoning. Meanwhile, knowledge carves nature at the joints, and in doing so, shows itself as discrete and as having definite properties. knowledge splits into diverse branches and subfields, forming a hierarchy just like ignorance, with subdivisions nearly as complex and differentiated as the universe that splits into such manifold fanning canopies of facts, events and phenomena. Where ignorance and knowledge differ in their typologies is that the tree of knowledge is richly populated with leaves, with nodes of positive content, while the tree of ignorance is bare as a skeleton. We can have physics or biological or chemical content knowledge as distinct branches on the tree of knowledge, but the came can't be said of ignorance.

It is not absolutely true that ignorance is always contentless or free of a relation to positive content. It is more like a tree in late autumn than one in the depths of winter, sparsely populated with leaves and entropic, shedding information states in a randomized pattern. For example in what I call positive ignorance such as as a false belief, conspiracy theory or delusion, there is an identifiable mental content which generates incorrect inferences which can be said to be caused by deeply set, stateful ignorance. That which causes, however, must have some effectual property, it must be a part of the causal history of the world. This content however is identified with the belief itself which can be said to have been "corrupted" by ignorance. It's not ignorance that has any sort of autonomous agency above and beyond or independent of the false belief's agency or causal powers. Ignorance rather expresses through the falsehood inherent to the misbelief which generates ignorant outcomes.

An entailment of the ignorance-first view is that knowledge can be incomplete both in the more obvious sense that all knowledge currently in existence does not

approach omniscience but also in the more nuanced sense that a given specific state of knowing can exist partially, that is, mixed with ignorance. How else can one person know more and another know less about a given subject, but still know something? Or how is it that one can claim probabilistic knowledge? What's more, knowledge of a fact does not entail a complete grasp of the fact and all its possible relations and combinations with the world. knowledge forms a system of relationships just as the world does. A typical formula in epistemic logic is that if one knows a proposition, and that proposition implies another, then one also knows the second proposition. This so called logical omniscience which has haunted epistemic logic evaporates once ignorance's role in structuring knowledge is properly appreciated.

So even in knowledge we do not escape ignorance. The incompleteness of knowledge both as we experiment and as the theory of knowledge, both turns on the fulcrum of inquiry. It is the fact of information incompleteness which compels the epistemic agent to close the gap between the gradient of external information source and their own deficient internal information state by seeking information to rectify its own awareness of that deficit. That is, to approach a state of equilibration between internal and external mutual information.

A major component to this interpretation of ignorance I wish to describe is also effectively a theory of thinking as information entropy reduction. In addition to examining ignorance, inquiry and their relationships with knowledge, there must be an idea of how cognition generally construed fits into this picture. It cannot be taken for an inscrutable black box. A naturalistic understanding of the cognitive psychology underlying these topics is essential. In essence we only think to the extent that we are ignorant, or rather to the extent that we are irritated by not knowing. All problem solving is a response to ignorance, otherwise the solution would be trivial, bordering on second nature or reflex. Knowing is to not have to think. It is the result of thinking, the resolution of doubt. Ignorance can be strongly correlated with the amount of randomness and unpredictability associated with a system of events, and we think so long as and until that disorder is processed and replaced by an intelligible pattern which then becomes part of a stable, static backdrop for actions and decisions.

The adoption of an ignorance-first viewpoint promises to productively update how problems in epistemology are conceived and interpreted. It courts the central problem of human existence. A world without ignorance or the perception thereof is a world without curiosity and wonder, without problems to solve, without purpose.

In the course of working out the meaning of this position we will come across many provocative and at times common sense defying observations. At the other side of this odyssey resides knowledge of what is arguably single most important cognitive reality—the thing that gives it its direction, and reason for existing: ignorance.

(Williamson, Timothy (10 October 2002). Knowledge and its Limits. Oxford University Press. doi:10.1093/019925656x.001.0001. ISBN 9780199256563.)

Preliminaries

Because ignorance is not only the simple lack of knowledge, there are several types of ignorance. It has a taxonomy. What follows will be much harder to understand unless the language employed is defined in advance. Given that ignorance is a prior, pervasive category and knowledge is a consequent, local, and relative category derived from it; given, that is, that knowledge emerges out of ignorance, it benefits a more adequate epistemology to distinguish the taxonomic subdivisions of ignorance. Some of these taxa are alluded to in colloquial speech, such as willful or blissful ignorance, however I view these as peripheral subspecies of more general taxa, namely what I call positive ignorance. Other classifications may sound foreign but once defined and described will become more recognizable.

There is a definitional problem associated with proposing a taxonomy of ignorance. Because it lacks tangible properties, there is noting concrete to compare it against. This absence of comparative anatomy however is an affliction that besets any hierarchy of abstractions and there is every reason to suppose ignorance can be fitted into differentiated logical types. The risk however is with

excessive variation and with the impossibility of precise scope and granularity in these types. Ignorance conceived in its totality is perfectly uniform and featureless, it is *unambiguously vague*. But various local instantiations of ignorance are more differentiated, but there is a problem of choice as to how to divide and distinguish them properly. The easiest way to solve this problem is to avoid it. I suggest that the taxonomy I propose is but a plus one, rather than a definitive, complete systematization.

The variety of ignorance types is a vital component of the theory of ignorance. The world of the unknown is more than mere absence but, when parsed, exhibits logical relations and structures that can be described as if they were objective realities. Ignorance is more than a state of mind or lack of a state of mind; in a manner of speaking, it is woven into the fabric of reality or a function of it. Little can be said about nothing. But if the unknown is actually a something, in terms of its practical, logical and cognitive consequences, then much can be said about it that has long been unsaid.

A note on terminology. Some effort was made to coin terms for the version of epistemology presented. “Inverted epistemology”, “anti-epistemology”, “entropic epistemology” are some that stood out. Though I consider it an imperfect name, I have simple settled on “epistemology”, occasionally referring to it as inverted, negative (neg-epistemic), or “outside-in” for reasons that will soon be explained.

It’s a curious bit of philosophical history that the words for the theory of knowledge and for the theory of ignorance were introduced into the English language simultaneously, but only the former caught on. Both words introduced by 19th century Scottish metaphysician James Frederick Ferrier. For knowledge he used the now standard term “epistemology.” The term he used for the theory of ignorance was the somewhat inelegant “agnoiology.“ Despite the brilliant ingenuity of this term and its obvious parallelism with the more conventionally used “epistemology”, it never quite stuck. The author admits to having agonized somewhat about whether to resurrect this term. The goal of this book after all is after all to reinvent and modernize Ferrier’s forgotten idea after centuries of shameful neglect.

Despite the appeal in reviving the dead word "agnoiology" I have decided to omit it from regular use. The reasoning for this is partly stylistic and partly a matter of a semantics. "Agnoiology" is frankly an ungainly word, and its foreignness to the typical reader might be off-putting. I also wish to put some distance between Ferrier and myself, as our accounts of ignorance while sharing several similarities differ in essential respects. While Ferrier gets the credit for originating the concept of a metaphysics of ignorance, I believe he does not go far enough, is wrong in several respects, and lacked a more comprehensive integration of his ideas with other relevant concepts.

(An account of Ferrier's lost theory of ignorance is included as an appendix. Because what I argue is not an extension of his ideas but a reinvention, it would be potentially misleading to accord it a more forward position in the flow of argument.)

I propose that the theory of ignorance is ultimately a branch of epistemology rather than adjacent to it. They share the same roots, project from the same trunk, split from the same bough. Or to rearrange the metaphor, the roots of the tree of knowledge grow from a soil of ignorance. I view ignorance and knowledge as inextricable facets of the same overall epistemic universe. So it makes more sense to label the theory of ignorance as an extension epistemology rather than its own field.

In recent years the term "agnotology" has been used. Agnotology however mainly finds employment within sociological and political science domains that address the cultural role and power dynamics of ignorance. The term does not designate a general theory of ignorance but rather the sociocultural impacts of it. Because we are primarily concerned with its more abstract theoretical structure, I will not use "agnotology" either to avoid confusing it with this more specific area of ignorance studies. That said, some attention will be given to the social and cultural role of ignorance as it is an important consequence of these lawful relationships.

I will sometimes also speak of ignorance and uncertainty, and knowledge and certainty respectively, somewhat interchangeably. While they are not quite the same things, they are strongly correlated. To the degree that knowledge is

uncertain it is probabilistic, it is contaminated with ignorance and shades into it. The degree of our certainty about a proposition quantifies the extent of our confidence in how ignorant versus knowledgeable we are about what it expresses. It may apply to both what we are ignorant of ("I am certain I don't know") and to what we know (I am uncertain I know that"), but the essence of uncertainty is ignorance and certainty knowledge. Uncertainty in particular is a central concept in the scheme of things, as it represents an important quantifiable construct which is the uniting thread that ties much together.

The terms "positive" and "negative" used to qualify knowledge and ignorance appear throughout. These are intended to be used in the sense of logical valance rather than emotional judgment, to refer to presence versus absence, active versus passive, assertion versus denial, rather than good versus bad. I also speak of "positivistic", or "propositional-assertional", knowledge to refer to the traditional view of knowledge in a way that may confuse those familiar with logical positivism. While logical positivism certainly resides squarely within this positive tradition, by "positivistic" I here mean knowledge-centric, a view of epistemology that interprets it in terms of overt, explicit epistemic propositions seamlessly united in a contiguous field of knowledge without discontinuous hiccups of ignorance interrupting its seamless structure. This is essentially the understanding of knowledge as ampliative inference. On the basis of what one already knows one derives new knowledge. knowledge systems are closed under their extension relations; they chain to other knowledge states with no reference to ignorance. This is the view of epistemology which sees it as entirely one sided, in terms of knowledge begetting knowledge. This ampliative perspective emphasizes that "most of us would prefer to "accentuate the positive" and focus attention on human abilities and powers, rather than disabilities and incapacities" and stress the knowable over the unknown as Rescher (2009, p ix) observers. As will be seen, the overemphasis on ampliative inference as a model for epistemology is wholly inadequate and inconsistent. Putting ignorance first, combined with various cutting edge theories from cognitive science including the predictive encoding paradigm, significantly clarifies the issue and evades the pitfalls of ampliative, positivistic epistemology.

Formal epistemic logic is particularly guilty of overemphasis on ampliative reasoning. Ignorance , if it is viewed at all, is interpreted as the negation of these positive propositional forms, the result of a mere application of the negation operator to these positive terms and nothing more. This assertional mode of knowledge views it as the manifest content of the mind and ignorance as essentially the negation of this manifest content. This positivistic epistemological outlook or bias is ignorance-blind and cannot see outside this notion of positive, usually conscious, knowledge and its extensions outward from its own states. This traditional viewpoint can be seen as an "inside-out" view, one that tries to build a theory of knowledge from positive propositional contents and chains of entailment from them. In this way, the inside-out view assumes what it is attempting to define. Meanwhile inverted epistemology seeks to develop an "outside-in" perspective. To derive, that is, knowledge from ignorance, and to understand them as constituting an irreducible reciprocal relationship. Given these points, ignorance first, outside-in epistemology is an inverted epistemology.

Finally while this treatment goes to extensive lengths to define ignorance it more or less assumes an intuitive meaning for knowledge. This is a little deadly, because ignorance cannot be built on ignorance, there has to be some knowledge of knowledge to orient around. While there is still much controversy about the definitional problem of knowledge and some opposition to the the popular definition of "justified true belief", I have no interest in getting sucked into that quagmire. For the current purposes, I will typically reference one of two definitions: one pragmatic and the other information-theoretic.

As for the pragmatic definition. Knowledge is that mental state which leads to effective action and the possession of which can be demonstrated through effective actions. Knowledge affords the possibilities that it knows. An agent can be said to know something if its mental state causes it to behave in such way that the predicted result encoded by the knowledge arises. The epistemic status of a mental state as knowledge is best judged by its consequences, by its material effects on the world at large and by the enhanced efficacy of the knower. Airplanes can be constructed, so it stands to reason that whatever actions and

thinking that went into building a flying plane was true and those who built it can justifiably say they know how to build it or understand the principles of flight.

This pragmatic definition of knowledge as proof by demonstration is admittedly imperfect, as one can arguably know something that cannot be acted out. And this is true of even simple pieces of knowledge such as knowledge of one's own memories, although depending on your definition of demonstration someone who knows their birthday might be able to show it by pointing to a birth certificate or by some other credible means. But it is this practical definition which is the most relevant from a cognitive-behavioral and scientific analysis. It is by far the most implicitly relied on measure of knowledgeability used in everyday life, from job interviews to academic exams, and it does most of what need of a definition without otherwise assuming an intuitive one.

Briefly to elaborate further on this knowledge by demonstration definition. It is important for knowledge to have some form of verification. Tying a knowledge claim to effective action provides a means to show how one knows. One does not strictly need to be able to demonstrate the fact itself, but can still demonstrate a justification for it, for example by providing a citation for data. I can still know about results from the Large Hadron Collider research facility without ever having stepped foot in it or flipping the on switch or even knowing how the machine works. This "proxy" demonstration, if acted out, would lead to tracing a chain of what has been called epistemic dependencies that would lead one to the facility. On one view failure to do so might show that one lacks the knowledge (Hardwig, 1985). Textbook knowledge might be of a weaker grade, because I cannot myself demonstrate its findings on the spot. This regression is eliminated however, if one posits that if the chain of dependencies would in principle terminate in some experiment that justifies the proposition, then the knowledge one has by proxy is valid. Ignorance is far simpler to prove. One can feign ignorance and play dumb, they can withhold knowledge but it is easier for someone to confirm that they don't know something than to prove they do. To prove I know Chinese I must pass a quiz. To prove I don't, flunking it will suffice. Because failure is more common than success, the odds of me flunking a test

correlating with not understanding the material is greater than serendipitously acing the test without knowledge of its subject matter.

Finally, anchoring knowledge in action permits the incorporation of exploratory and investigational behavior into epistemology, the key to understanding inquiry as the pursuit of ignorance resolution. When you don't know something but want to find out you are inclined to investigate, to take exploratory actions that mitigate your ignorance. Ignorance is particular in this regard because the only recourse in response to its indubitable presence is information seeking activity through inquiry. One might know without necessarily having to take action, but all inquiry, all goal-driven processes of ignorance remediation, necessarily entails doing something about it.

Turning now to the informational definition of knowledge, first let's make it clear that this comparison does not equate them. Information and knowledge are interestingly related, but they are not substitutable. Instead this second definition is more like one for knowledge *formation* than knowledge representation. The ins and outs of this will all be explained in overwhelming detail in due course. A brief, high level description of here will suffice. In the plainest terms knowledge is the signal extracted from the noise. The initially uninterpreted object of potential knowledge is a swirl of alternatives each without any weighted estimations as to which is more likely to be the case. Through some otherwise unspecified process of neural entrainment, information streams observed from the object eventually collapse into a pattern of perceptual invariance. Weighted estimation values are assigned to the various alternatives, which "collapses a pocket" that then accelerates into self-organizing until a relatively stable, controlled model of the object's behavior and properties is constructed. A knowledge representation is therefore like a kind of residual concretion left over from the exclusion of irrelevant details. It is, as we say, what happens when you "boil it down." (This view is inspired by the so called neural-constructivist paradigm of learning)

Together these two definitions should cover enough ground between them to see the theory through.

With these provisos out of the way we now move onto the task of describing

an overview of the type system for ignorance. The following section can effectively be read as a synopsis. It lists all the component parts which the rest of the book assembles.

Provided is a thumbnail outline and high level mind map of the conceptual taxonomy of ignorance and terminology that will be used. This is intended to be a reference guide to help keep track of the various terms and relationships that inform the bulk of the story and to preempt any ambiguity that might arise from using these terms. I include it here rather than as a backended supplement or glossary because I believe it will greatly assist in the reader's comprehension to put it up front. The categories and terms, though listed top to bottom and in sequence, are not meant to be independent of one another and the order in which they are listed does not indicate descending degrees of conceptual importance from top to bottom. Instead, concepts are organized in groups in terms of their relevance to each other. Neither is the list necessarily exhaustive. It's intended only to provide an overall breakdown of the major conceptual distinctions that will inform the discussion. For expedience, much will simply be asserted here and supported later in the main body of the text.

Absolute ignorance

Definition: Ignorance with a capital I. The sum total of all undifferentiated ignorance. Everything that everybody doesn't know. All unknown unknowns fused into one. Including what is putatively unknowable and all that is contingently unknown to all knowers. Ignorance as a type rather than a token. Ignorance taken as a monadic, perfectly uniform whole.

Physical ignorance:

Definition: Ignorance which is embedded or deeply correlated to physical objects, relations, and processes. These include limitations on information access prescribed by the laws of physics such as the quantum uncertainty principle, black holes, entropy, temporal relations, speed of light.

Subcategories:

— *Temporal constraints*: what is unknowable is connected to our relative standpoint in time. The knowability of past and future is asymmetric and

generally disfavors knowledge of the future, but neither are as knowable as the present. Generally the further away in time an event is from the present, the harder it is to know anything about it, because in the case of the past, information about the fact has decayed through entropic dissipation for longer. In the case of the future events, for all practical purposes they are mere notional possibilities and do not yet exist. Making matters worse the laws of physics prohibit information from being communicated from the future to the past. Though there may be some retrocausal quantum weirdness that is the exception to the rule, from the standpoint of the typical observer this makes no practical difference. Regardless of how you slice it, any given observer has a temporally constrained incomplete information state with respect to past, present, and future. Temporal ignorance is of great explanatory importance and situates many of the other problems in their proper context, notably in relation to entropy.

— *Irreversible entropic processes*: thermal entropy increases the amount of randomness and disorder in a system and hence complicates the requirements for an adequate description—hence knowledge—of the system. Increased disorder entails increased uncertainty and irregularity about the disposition of events. Information that undergoes entropic decay becomes irrecoverable and passes into a "high ignorance" state. While congruent with information (Shannon) entropy, thermodynamic entropy is not the same and has a distinct explanatory role apart from information entropy. In particular, because entropy is irreversible, information about the previous state of the world is always lost. This fact roots temporal ignorance.

—*Physical information bounds*: phenomena such as black holes, the speed of light, light cones, the quantum uncertainty principle and thermal entropy which firmly restrict, funnel or compress access to information. Many of these phenomena are of small practical significance in day to day life but they are of great theoretical importance for establishing the premise that ignorance in a manner of speaking is actually embedded into the objective world and is not merely a relation of the mind or lack thereof to the world.

Computational limits

If the brain is to be considered an information processor then various

theorems about formal undecidability and incompleteness ought be applicable to it or at least helpful in determining the upper limits of what minds can know. The issue of whether the brain is some kind of computer is hotly debated but this is not the subject at hand. Of paramount interest is how limitations on how axiomatic formalisms might speak to limitations about the mind. This does not absolutely require the the mind be reducible to a formal system. What it does require is a point of contact in which the human mind touches the meaning of these proofs, having grown within itself these theorems about the limitations of what can be done with some of the most powerful tools of thought such as mathematics and computation.

Subcategories:

— *Undecidable propositions: *certain propositions that have no answer criteria that are formally satisfiable. There are no grounds that are both consistent and complete for deciding whether the proposition is definitely true or false, though it may very well be either. An undecidable proposition is not a *truthless* proposition, it is one whose truth cannot be proven.

- *Decidably undecidable propositions*: A neologistic phrase that designates any proposition which can be determined to be undecidable by some proof mechanism. For which, that is, a second proposition referencing it and asking whether it is undecidable can be answered with a definite yes. Turing for instance decidedly proved that the halting problem is undecidable. Gödel proved through his incompleteness theorems certain fundamental restrictions of axiomatic systems. These propositions are equivalent to a *knowably unknowable unknown*.

- *Incompleteness*: Fundamental restrictions on the provability of axiomatic systems. Formal systems must trade off between completeness and consistency, meaning between containing a proof for every true statement and the perfect absence of contradictions. See also *information incompleteness*

- *Constraints on memory and processing*: The problem of information overabundance in relation to the finitude of computational resources. The brain or any computational system has a theoretical upper limit on how much memory it can store and a fixed processing speed. Unless that first limit is larger than the

descriptive complexity of the universe that contains it or the second limit can process information faster than the speed of light, both obvious absurdities, any given mind must be doomed to some level of ignorance about the universe by quantitative necessity. Any number of facts that requires more cognitive resources than those constraints allow must be unknowable, at least directly by a single individual, because the amount of resources required to process it exceeds the neuronal channel capacity. There is more out there than can be computed by any inference engine, including human brains, and it is unclear that the brain is able to compactly "minify" such a massive abundance of information efficiently in compact generalizations that capture the properties in a prototypal abstraction. Most notably such minified abstractions are natural laws. Natural law is a great exception however. More likely it just has to be selective about what it chooses to know. In fact somewhat ironically there is more *knowledge* than one can know, because at a given time one ultimately has to sacrifice the opportunity to know one thing rather than another. No one can read every book ever written, not because their life is too short, but because there are strict physiological limits to how fast they can read.

Cognitive closure:

The thesis that there are certain unthinkable thoughts and unknowable unknowns owing to the finite, determinate and self-contained nature of the brain or mind. Somewhat related to constraints on memory and processing, except cognitive closure does not have to do with how much or how fast things can be known, but rather with what the brain is intrinsically capable of understanding. The school of thought called mysterianism for instance holds that the so called hard problem of consciousness is unsolvable by humans. According to them it is humanly impossible to understand how conscious states manifest out of the activity of brain tissues.

Cognitive closure proposes that we can only think the kinds of thoughts our cognitive architecture permits, and everything outside this bound is lost to absolute ignorance. It is to be argued that the question of cognitive closure is itself cognitively closed to us, since, as Wittgenstein stated in the *Tractatus*, "for in order to be able to draw a limit to thought, we should have to find both sides

of the limit thinkable (i.e we should have to be able to think what cannot be thought)." The thesis of cognitive closure is therefore undecidable, and decidably so, and therefore closed. To prove it, we would have to find evidence of a thought which can't be thought, an obvious absurdity. At the same time, no matter how many thoughts we do have, there is no way to infer there are concepts outside the bounds of what we have thought. Perhaps there are alien thoughts in alien brains that represent everything differently. This invites various vexing paradoxes, because the limit of thought straddles both sides of what is thinkable without belonging to either category, caught in a state of indetermination between true and false. Or the limit does not exist, but we would never be able to prove it did or didn't regardless.

The extent to which closure holds true beyond the one limiting case of itself is necessarily ambiguous and itself unknowable, because it is impossible to fathom the exact scope and extent of what we can't conceive of without conceiving of it. Whether cognitive closure exists can't be ruled out and, as will be demonstrated, is an undecidable proposition for which neither justifications for affirming or denying can be asserted.

Subcategories:

- The self-referential paradox of cognitive closure itself, the possible inability to know it's true that there are thoughts we can't think.

- Possibly the mystery of consciousness

II. Relative Ignorance

Definition: Ignorance which is relative to a particular agent, circumstance, condition, belief state, cultures, upbringing, history, or inequality between two sides of a gradient. Depending on the facts at hand it may or may not be remediable. This is ignorance that is potentially discoverable and learnable and can be converted into knowledge.

a.) *Inaccessible contingent information*: facts which are knowable in principle but for which there is no chain of events connecting an agent to it through causal acquaintance. The laws of physics do not rule out knowing these things. There is nothing special about these facts that necessitates they should be unknown. What distinguishes them is simply that there is no way for a given

agent to acquaint themselves with it at a given moment by establishing a causal link to it. To know is to be caused to know by the object of knowledge (Goldman, 1967). Such restrictions are relative because another agent may have an have access to the facts because the worldline of their causal history is better positioned to intersect with it.

c.) *Knowledge gradients*: the uneven distribution of positive information and negative ignorance between subjects or between a subject and information environment or source more generally. A teacher knows more about the topic than the student; this imbalance of knowledge and ignorance between them is the reason why the student is in school. knowledge flows from the position of the higher concentration gradient to the lower gradient in a manner loosely similar to thermal dissipation. As one learns from a given information source, the gradient approaches equilibrium. While it is likely rare for true equivalence to be reached even in cases of mastery and expertise, and this information flow is not so simple as the transfer of energy, the approach towards this ideal limit continues until the learner's goals have been satisfied or the process is otherwise interrupted. knowledge gradients are important for understanding the general epistemic infrastructure of society and how knowledge and information circulates within it by tending to go from a denser higher concentration toward the vacant lower concentration "heat sink" of negative ignorance. The analogy to a thermal system is imperfect, given that communicated knowledge does not have to be subtracted from the higher gradient to add it to the lower in line with the law of energy conservation, but is instead preserved across both sides of the gradient (unlike heat energy for which the exchange is exclusive: an increase in one area necessarily implies a decrease in another). Additionally someone who is ignorant but wishes to correct it actively "moves against" this gradient rather than passively awaits information to come to them. In general however the same principle of a differential slope and the tendency for flows to go from higher to lower concentrations along it holds in the thermal and informational cases, and this analogy does much to clarify the relationship between knowledge and ignorance.

d.) *Information incompleteness*: the imperfect condition of a cognitive agent

relative to its environment, which puts it at a continual information deficit and inequality with respect to the larger system that envelopes it, and which serves as the motive principle for learning. Learning is the effort to reduce that discrepancy by increasing the mutual information shared between the internal and external agent/world boundary. Information incompleteness can be thought of as the inverse of a knowledge gradient. It is ignorance that draws in knowledge by "expending energy" through effortful learning to move against the gradient to close the gap and lessen its imperfection. Examples: minds, brains and sensory perception, exploratory behavior, living things in relation to their environment in general, science, police detective investigations, any intentional learning effort. ("There are more things in Heaven and Earth, Horatio, than are dreamt of in your philosophy.")

Terminology:

- *Uncertainty*: a state of indeterminacy as to the outcome of an event, the truth status of a proposition, or the result of a decision. Typically defined in elementary terms of bits. Implicit to the concept of uncertainty is the recognition of alternative "what if" possibilities that have not been disambiguated or excluded from consideration. Approximately, the amount of information an agent gains is inversely proportional to the amount of uncertainty it perceives (Kullback, S.; Leibler, R.A. (1951). "On information and sufficiency". Annals of Mathematical Statistics. 22 (1): 79–86. doi:10.1214/aoms/1177729694. JSTOR 2236703. MR 0039968.). Information reduces uncertainty by chopping down excess possibilities, by reducing Shannon entropy, by answering questions. Uncertainty is deeply interwoven with ignorance although they are distinguishable. (The same goes for certainty and knowledge.) In this work uncertainty is treated as a fundamental primitive concept. See also *information entropy*

- *Information (Shannon) entropy*: The informational analog of thermal entropy. This work assumes the conventional definition of Shannon entropy used in information theory. Info-entropy is the measure of the noise or vagueness in a message. Information is any action, signal or stimulus that reduces uncertainty. Hence the information entropy of a variable is the average amount of

randomness represented by its possible values, its degree of ambiguity. The information necessary to characterize a system is equal to the amount of bits necessary to satisfy the amount of uncertainty proportional to the entropy. Each bit has an uncertainty or entropy of .5 because it may occupy one of two possible exclusive binary states, either 0 or 1 (yes or no, true or false, on or off, absent or present, here or there, etc). It is the measure of how much surprise a message contains, that is, how improbable a given token within an ensemble is relative to the alternatives. Improbable events yield much information, while probable events less. Information entropy can also be thought of as the measure of the diversity or irregularity inherent in a signal, message, process or really anything. Given that probability is a measure of ignorance, information entropy is a helpful concept for analyzing ignorance in precise terms. The more randomness and irregularity in a system, the greater the *Kolmogorov (descriptive) complexity* required to specify the shortest possible description of a program needed to generate its behavior. Epistemologically this is another way of saying the more irregular, unusual, and complex a problem, the more likely we are to be ignorant about it because of the more work we need to do to codify it into a knowledge representation.

- *Negentropy*: Formally defined, information negentropy (negative entropy) is the difference between the current state of entropy (randomness, disorder) of a probability distribution and the maximally disordered entropy state. The values of negentropy and entropy are inversely proportional. The higher the negentropy the more ordered or predictable the system, the further away from a normal distribution it becomes, which means the further away from totally randomized ambiguity and vice versa. Negentropic processes, insofar as they are local and temporary, do not contradict the second law of thermodynamics which states that entropy always increases. Life itself is negentropic. More to the point, *thinking* is negentropic. This is not an alien idea, indeed, there is a rich literature on the topic of how cognitive structure self-organizes through negentropic processes ([https://www.researchgate.net/profile/Damian-Kelty-Stephen/publication/40455355_Dynamics_of_Representational_Change_Entropy_Action_an

d_Cognition/links/0912f50bccf485bc58000000/Dynamics-of-Representational-Change-Entropy-Action-and-Cognition.pdf]) One of the driving objectives of this work is to build on this literature, specifically to synthesize such notions into an epistemology of ignorance. The quantity of entropy in a system is directly correlated with how much (negative) ignorance, or how many opportunities for ignorance are present within the distribution. To form a mental representation is non-superficially analogous to the crystallization of an inorganic solid. knowledge on one level is the detection of invariance. It is a stable and reliable consolidation of possibilities which are initially unpredictable but then restricted to comfortably expected bounds.

- *inside-out/outside-in epistemology*: The typical view of knowledge that has been its intuitive working model for millennia is that knowledge begets knowledge. Knowledge is an extendable system of rules, states, and operations nothing outside of which is relevant to its logic. It is grounded in ampliative inference, with known truths extending to other unknown truths. It begins with some foundational premise, some irreducible axiom or initial knowledge state. Then through the application of inference rules to operation of reasoning otherwise unspecified, it extends to more knowledge. Epistemology on this view is "inside-out." It assumes that derived knowledge is linked to prior knowledge through unbroken inferential chains and that it expands outward as knowledge grows. This assumes that the analysis in question has any appreciation for ignorance whatsoever, that it even recognizes a boundary, an outside. Some formalisms like epistemic logic assume epistemically ideal conditions that recognize no such divide. Inverted epistemology in contrast, is "outside-in." It depicts ignorance as the starting point for knowledge. knowledge is drawn toward it in the effort to collapse the gap between it and to seal the envelope. knowledge is worked out from ignorance, and ignorance tapers off into knowledge through processes of learning and inquiry. Ignorance, looked at this one angle, begets knowledge. Ignorance in this view is ripe with potentialities and counterfactual alternatives, rather than an inert and infertile nothingness.

- *Negative (simple) ignorance*: the conventional dictionary definition of ignorance as the lack of knowledge. This form of ignorance is contentless, void,

and defined strictly as the absence or negation of true knowledge. It can also be used, somewhat interchangeably but with technical distinctions, to mean a lack of information. Uncertainty is an important type of negative ignorance, because it is quantifiable on an elementary level as a bit or a probability expectation value as a number between 0 and 1. If an epistemic possibility were a bit, negative ignorance would be the 0 to positive knowledge's 1.

- *Positive ignorance*: a type of ignorance that has or is inextricably associated with and pervades mental representations. These are mental states that are distortions of reality. This is when one is wrong, believes incorrectly and is unwilling or unable to realize it, rather than merely the lack of true belief. Such ignorance is distorting: it falsely represents states of affairs rather than merely lacks a true representation. An example of this would be the unwitting belief in a falsehood, superstition, fallacies, self-deceptions, delusions, hallucinations, various mental illnesses, cognitive bias or a false memory. It is not enough to say that one merely doesn't know the truth in these cases. Instead, mental representations project a version of events which is incongruent with reality. Such false beliefs do not simply lack the truth, they deface it. When positive ignorance is in play, the person is reacting to something rooted in their faulty reasoning about the world.

- *Complex ignorance*: A system of ignorance that is elaborate and built up. Complex ignorance is organized and systematic. It has a logic, albeit a false one, and can extend inferences that draw spurious conclusions which nevertheless sustain the system's development and do not contradict its rules. Such a complex often superimposes on a belief system and consists of a combination of mutually supporting positive and negative ignorance that scaffold each other. It may even contain elements of knowledge misappropriated to enhance the overall false design. Complex ignorance can either afflict an agent or be imposed on another as in an adversarial plot or stratagem. It may result in wrongly invested intelligence rather than a complete absence of it, or can lead to over-engineering, over-planning or unintended consequences and Hindenburg disasters. These are networks of misbeliefs, rationalizations, trojan horses, conspiracy theories, cult doctrines, political disinformation and campaigns, moral panics, extremist

ideologies, and pseudoscience. Perhaps even a dream when someone does not know they are dreaming. Complex ignorance often masks itself in the guise of certitude and knowledge. It actively misdirects and misleads thought to reliably produce flawed conclusions and drive one deeper into ignorance. It is not a single corrupted belief, but a method for organizing beliefs falsely. It is a pattern of unwitting falsehoods, strategic ignorance, ideological ignorance. Complex ignorance must have some positive content, negative ignorance only plays a supporting or maintaining role for example by excluding or ignoring truths inconsistent with the framework.

- *Structural (causal) ignorance*: ignorance which is causal, to which the outcome of a chain of events is attributable . There is nothing mystical about the idea. As is plainly obvious, what we don't know still plays a role in the unfolding of events, the causal nexus does not stop at what we are aware of. Structural ignorance is that which organizes available, remaining epistemic options after one has come to your "wits end." Thinking that falls short of a complete analysis of the problem, must then work within the context of ignorance associated with what remains unknown about the problem. This is ignorance that specifies what is left over when knowing stops. It can be visualized as funneling or channeling. Structural ignorance is usually to be found when the exclusion of alternatives determines the result of conduct and the choice between recognizable alternatives. One makes a decision not always because it is the best one, but because other options that might be better are unknown. The fact these other alternatives are unknown shapes the decision space. The absence of other options leaves open the only remaining choice. Structural ignorance is what gives us what's left over when everything else has been excluded, cut short, made inaccessible, clouded over with ambiguity or that slipped by unnoticed. It can be useful, as with the pruning of a search tree saving processing costs by cutting out inferably irrelevant branches rather than executing a thorough brute force search method that checks every value inefficiently. Or it can be useful to help concentrate and regulate attention by selectively ignoring irrelevant details. Ideas and projects are often the result of tactical decisions to choose not to go down particular paths of wasteful thought. It is to ignore, to exclude or be excluded.

Selective pruning which provides the eventual shape of the output just as much as the inclusion of contents ultimately determines the outcome. Inclusion and exclusion are not the same process expressed differently but two functionally cooperating processes that determine the result by moving in opposite directions, reducing possibilities in different ways and converging at the same set point like burning a candle at both ends. To ignore and to pay attention are both methods which include in one set and exclude in another. The very fact that some things were unknown either by choice or accident, is what explains the outcome. Indeed too much knowledge or awareness of alternatives can be inhibiting, leading to so called “analysis paralysis” or “information overload” in which time is wasted and action withheld dwelling on all the options. Structural ignorance is when the progression of a process is a consequence of information deprivation. Examples include trial and error problem solving and process of elimination, synaptic pruning in brain development, habits, probabilistic reasoning, evolution by natural selection, senescence, satisficing, various forms of conflict, and many more phenomena. See also *agniogenesis*

- *Potential knowledge*: Often we speak of ignorance as though it were merely knowledge we could possess if conditions were met. In this respect, there is a supposed logical homology between ignorance and knowledge. Ignorance is merely knowledge which is inert, lacking the cue to activate it. It is possible knowledge. Ignorance is alethic, modal, in nature and what is unknown is knowledge waiting to be discovered. This interpretation of ignorance is flawed however but still useful particularly when talking about learning, discovery and inquiry processes. It is flawed because the isomorphism that it asserts holds between knowledge and ignorance as potential knowledge is inaccurate. knowledge is representational, while in a latent stage what one comes to learn is an unstructured “bitstream” of information signals lacking in any coherent unity in in of itself. Knowledge lends unity and coherence to a correlated cluster of information signals, that identity is not intrinsic to the information stream. Knowledge is not acquired, only information is. Knowledge is constructed out of or in response to information but it is false to assume what is to be learned is analogous to what is learned. If ignorance were always potential knowledge, by

this same token it would mean something like Plato's concept of knowledge as anamnesis, knowledge as remembering or unforgetting. The only potential which is activated is already present in the brain, not the external information signal that prompts its activation and which we think of as the cause of learning. Lastly there may be unknowns for which there is no possibility of knowing for one reason or another.

- *Negepistemic (negative epistemic)*:. While I wish to spare the reader from inflicting too many neologisms on them, the fact of the matter is there are few terms to explain ignorance-knowledge relationships. Specifically there is no word to describe situations which ignorance and knowledge are simultaneously equally relevant, or when the relationship between both needs to be referred to rather than either in separation. For example as when referring to a known unknown. Or when thinking about a knowledge gradient, the difference between what you know and what I know. So I propose the term negepistemic, a contraction of negative and epistemic, to refer to these use cases. It is not the most elegant word, but it saves quite bit of breath. And it is better than the second option I might have gone with, the simply ghastly "agnioepistemic." See also: *negentropy*

- *Negepistemic ordered pairs*: These are known-knowns, known-unknowns, and unknowns unknowns (and debatably the edge case unknown knowns). Inspired by a passing remark by former US Secretary of Defense Donald Rumsfeld, these pairs of terms have become something of a shorthand notation for ignorance-knowledge relations in the literature. These pairs are an extremely handy and succinct means to express the every meaningful permutation of knowledge and ignorance relational pairings. While researching the history of this shorthand, I noticed that there was no term to refer to the concept of these pairings as a whole. It is such a concise and compact way to express negepistemic relations that it would be a shame for it to go unlabeled. This is why I propose the term negepistemic ordered pairs. They are negepistemic in the sense explained above, and they fit the definition of an ordered pair, a collection of two terms (x , y) the order of which is meaningful unless x = y.

- *Modal negepistemic tuples*: same as the above but with modal semantics

and some twists and additions. Because some of the combinations have two terms and others three, they are called tuples, (technically n-tuples) which means an ordered list of elements of finite but arbitrary size. The complexity of the language is unfortunately increased, and the number of permutations of terms doubled, by introducing modal semantics. So it is understandable if these distinctions are hard to follow and parse at first. The terms are: knowable knowns, knowably knowable knowns, knowable unknowns, knowably unknowable unknowns; unknowable unknowns, unknowably unknowable unknowns , unknowably knowable knowns, and unknowably unknowable knowns. At a glance it is not exactly clear what these terms refer to, but each has referents and are greatly clarified by example. Some, like a knowable known are tautological edge cases with less meaning than others but they are included in order to enumerate every combination. Rather than explain these here I explain them in chapter 7 section 2. These modal concepts are more complicated and ungainly than the basic pairs, but are helpful in the analysis of the conditions of epistemic possibility.

- *Meta-ignorance* (also self-ignorance, ignorance blindness, ignoration): the ignorance of ignorance. This is second order ignorance, for which no ready-to-hand knowledge can provide inroads into. Unlike absolute ignorance, it is particular to a given subject's self-awareness of their knowledge and its limits. It is a type of unknown-unknown, except it reflects the agent's relation to their own ignorance rather than to an external affair or fact about the world. It is characterized by a blindspot in cognitive self-appraisal and inventorying which fails to register and pick up unknowns and tag them as known unknowns. Said unknowns slip through the cracks of the self-reflective dragnet used to introspect on one's knowledge. Often this ignorance is only remediable by lucky accident or chance exposure, or by someone else altruistically noticing that you don't know what they know and informing you. Or else it is fixed by a process of reasoning which conducts a series of inference steps that builds up to a new insight. It can also mean the lack of insight into the concept of ignorance rather than distinct instances. For example, traditional epistemology's lack of appreciation for ignorance as a pertinent subject.

- *Framing ignorance*: Ignorance which is caused by the overemphasis and overestimation of certain knowledge. Similar but different to the cognitive bias called functional fixedness, which is the tendency to rigidly use an object only in the context in which it is usually used. Assumptions determine what is reasoned about. Or valid data which is used inappropriately because of false assumptions or bias. It can be an overgeneralization that attempts to interpret what would otherwise be part of the truth as the whole truth. It can be a stereotype, a hasty generalization or a spurious correlation that magnifies a grain of truth into an oversized glacier. The point is, the inflated proportions of a partial truth or the appearance of truth does not discredit the underlying fragmentary truth, but only cast doubt on the method by which it was generalized. This is a kind of misappropriation of knowledge in which one chooses to interpret a problem using an inadequate framework or by misusing a valid framework. This form of ignorance is notable for involving the misapplication or overgeneralization of knowledge rather than the lack of it. It is different from complex ignorance because it refers to the misuse of actual knowledge, although it may comprise part of a complex. Examples include certain cognitive biases such as cherry picking, functional fixedness, cognitive processes of priming and anchoring, greedy reductionism, selective memory. "When all you have is a hammer everything looks like a nail."

- *Agniogenesis*: A neologism denoting a process that generates ignorance. A subtype of causal ignorance, agniogenesis is the increased diffusion of information entropy into the environment, most often by the actions of misdirected or ill-intentioned cognitive agents. It is that which "muddies the waters," pumps more confusion into the world. A lie is a good example. A lie is a linguistic dart filled with the tranquilizing poison of ignorance shot by one person into another. A lie produces positive ignorance in another person's mind, which the liar then uses to structure and scaffold their schemes by working with and around what the other person now falsely believes and doesn't know. State propaganda manufactures ignorance on an industrial scale. Forged currency, fools gold, camouflage, scams, fraud, high tech "deep fakes" are all examples. Agniogenesis can be either generative (positive) in the sense described above or

destructive (negative). It can manifest also as a simple act of destruction-of evidence, sources of information, encoded memory, or knowledge bases. As in the death of a renowned scholar, censorship, the burning of the Library of Alexandria, a hard drive wipe, or Alzheimer's disease. With the passing of each of these events the possibility of knowing or recalling something is eliminated. Agniogenesis differs from plain entropy in that it is intentional or agentive. It is either caused by or affects epistemic agents and is not an otherwise indifferent natural tendency.

The Logical Form of Ignorance

> …we are often unaware of the scope and structure of our ignorance. Ignorance is not just a blank space on a person's mental map. It has contours and coherence, and for all I know rules of operation as well - Thomas Pynchon

Ignorance is fundamental. It is a function of the causal structure of things—i.e., how various causal histories impact what can be discovered by a cognitive agent at a given moment in time. It also has a taxonomy, it is possible to identify various grades, types, and expressions of ignorance.

The ways ignorance impacts human life are practically inexhaustible. The ethics and politics of ignorance, willful ignorance; its social repercussions, its implications in business and relationships— the experience of it from a mortal perspective in all its iterations. Our primary concern here however is with the theoretical structure its logic at the root level that remains when you strip away all contingent or specific details. In a word, the objective is to identify the *Laws of Ignorance*. Ideally this approach would allow us to think about ignorance in a domain-general sense, not just in terms of how humans experience it but how *any possible cognitive agent* might. That includes at the lower extremes any lifeforms with even rudimentary capacities for information processing, like protozoa or the eukaryotic cells of an animal, though the simpler the organism the simpler its relationship to ignorance and the less it can be characterized in terms of knowledge representations and the more in terms of rudimentary

information signals. At the higher range, it includes people as the most familiar subject of ignorance, or any sapient beings and possibly artificial intelligences.

In the course of this examination of the logic of ignorance it is also expected that an adequate description of ignorance provides the natural backdrop and context in which knowledge unfolds. Short of some hypothetical omniscience, knowledge is always in some way related to ignorance, which is to say that no real-world epistemic agent operates in an environment completely devoid of unknowns and it within and between this outer backdrop that knowledge is to be properly situated.

A central pillar to this theory of ignorance is an externalist variant of epistemological realism. Namely, it's the claim that uncertainty is an objective feature of the environment over and above a state of subjective cognitive indecision. At first this may sound like a presupposition that hinges on an unknown premise that the universe is innately indeterministic. Such a position has its precedents and been defended before. C.S Peirce is among the most prominent with his idea of tychism, the view that the universe is ruled by absolute chance and that it is an objective property of nature.. According to Peirce, natural laws are not really what structure the physical world from the top down. The regularities of the universe are like habits which have "congealed" from an overarching tendency towards chance and randomness. It is these statistically induced "habits" of the natural world, and the causes that generate them, which are labeled as natural laws. Uncertainty, furthermore claims Peirce, is not just an irritant or defect of the mind, but an important aspect of its function. Without such fruitful uncertainty, the mind would lose a certain slack or flexibility, thus becoming rigidified and senescent. No mental activity, writes Peirce, "seems to be necessary or inevitable in its character. In whatever manner the mind has reacted under a given sensation, it is the more likely to react again; were this, however, an absolute necessity, habits would become wooden and ineradicable, and, no room being left for the formation of new habits, intellectual life would come to a speedy close. Thus, the uncertainty of the mental law is no mere defect of it, but on the contrary of its essence. (Peirce, The Law of Mind). Such negative capability as it perhaps turns out is far more important for how the

mind operates than might at first be suspected. On Peirce's view, the mind is an organ that is sensitive to chance and quickly adapts to acquire new habits which help stabilize by anchoring it to islands of regularity in an otherwise chaotic universe.

Treating uncertainty as fuel for the mind that is "out there" has several payoffs. Numerous areas of science and math, from information theory to quantum and statistical mechanics to probability theory , all assume uncertainty or indeterminism as properties of the system described. Or at very least, they assume that the system cannot be adequately described without resorting to these probabilistic methods. Skepticism about this uncertainty actually being an attribute of the system rather than a confusion of the limitations of the mind in relation to it, the so called mind-projection fallacy, is to be addressed in turn. Suffice it to say here that it is really how the mind relates to uncertainty that is of utmost explanatory value.

Pivotal support for this thesis is to be found in the pragmatistic argument that, were the quantities defined in these fields imaginary, it is not clear what is actually being measured or why they should be so effective. Why should doubt be so quantifiable? . The more this question is is pursued to its conclusion, the more it becomes clear that ignorance and uncertainty is an attribute of how the world is constituted, it is what the mind "works with" to obtain knowledge.

This approach to ignorance isn't to insinuate that it can be understood without reference to minds. Rather, it is best looked as the product of agent-world interactions.

Usually the more scientifically precise attempts to analyze epistemic concerns are, the more the conversation tends to become less about high level hierarchical knowledge representations and more about low level information states. There are no precise means to describe the brain as a knowledge-system or representational engine, but great strides have been made describing it as an information processor. Thanks in large part to David Marr's levels of description approach, there has been much progress in studying brain circuitry as computational systems. Comparatively it is more difficult to pinpoint how exactly knowledge is instantiated in the brain. This suggests that there is a

difference between computational competence and deeper epistemic comprehension. Systems lacking in the capacity for higher representations lack knowledge, but they still may exhibit sensors that allow them to extract information signals from noisy environments and basic memory storage mechanisms to persist that information for future use. One typically attributes information processing to computers, but it is more dubious to attribute knowledge to them. In conventional parlance it is said computers have competence but lack comprehension. Uncertainty is still relevant for these computational systems in so far as they are information processors. Such rudimentary information processing, I'd argue, is the minimally sufficient conditions for which one can meaningfully describe the behavior of a system as conditioned in terms of ignorance. In other words ignorance can be taken a step down in level of description to apply to the informational level, whereas knowledge cannot be so reduced. The reason that ignorance can be analyzed at a more primitive level is because its close association with uncertainty.

You can't have knowledge without information, although you can have information without knowledge. Given that information is itself a measure of uncertainty it then follows that you can still speak of ignorance in systems that are otherwise lacking in knowledge insofar as they process information. Much hinges on this key point. Because modern physics, as will be detailed in Part II, works in large part with an information-based ontology, the fact that both mind and world are analyzable in terms of information allows them to be subject to the same framework of uncertainty.

The distinction between knowledge and information, or more accurately information and ignorance, given that information is a measure of uncertainty rather than certainty, is not an absolute and sharp line of demarcation. The view supported here is that knowledge organizes information, associating related patterns of information into a structured representation. Information without knowledge is of far more limited utility because the information by itself does not tell you what to do with it. knowledge is what tells you what to do with information.

Uncertainty is more fundamental than certainty because it is prior to all

understanding. It is there before one has an idea of the facts and has made up one's mind, before any thinking happens. And while it is possible to exclude wrong hypotheses, it is harder to know that all the data one knows has exhausted all the potential for discovery and perfect certainty thereby achieved. This fallibilism is a necessary entailment of the ignorance-first view. It is easier to measure what to exclude, but harder to judge the completeness of the knowledge which remains after ruling out falsehoods . The foremost reason for the undecidability of perfect certainty is related to Meno's paradox: if one knows what one is inquiring about, then there is no need to investigate. If one doesn't know, then investigation is impossible. The larger point is that it is not always realistic to know where to draw the boundary between what one seeks to determine and the open expanse of everything else that could be relevant or is irrelevant. For it to be otherwise one must presume some initial understanding of the object of investigation, but in a state of total ignorance this boundary cannot be fixed. Neither is learning what is unknown like a boiling teakettle that will whistle once all has been learned; there may not be a clear and definite signal from outside indicating that the subject has been exhausted of all potential findings. Such is the nature of induction that it never arrives at perfect completeness. Writing against the existence of absolute necessity—the idea that all natural events must be determined by natural law— Peirce notes that all empirical inferences are provisional acts of approximation indefinitely suspended from reaching a terminal state of totally complete information.

> By saying that we draw the inference *provisionally,* I mean that we do not hold that we have reached any assigned degree of approximation as yet, but only hold that if your experience be indefinitely extended… then our approximation will become indefinitely close in the long run; that is to say, close to the experience *to come* (not merely close by the exhaustion of a finite collection) (Peirce, 1892, p 327)

From an ignorance first perspective, knowledge can be conceived as ignorance subtraction, If additive knowledge is to map the known to the unknown through ampliative inferences such as analogy or induction, what

might be called subtractive knowledge goes from the unknown to the known, through a process of reductive exclusion of information free parameters. It is what could be termed *reductive* inference. Differently put, this act of reduction is the process by which knowledge is gained by giving one less irrelevant things to know, as when a biologist fixates on describing life in terms of organisms, organ systems, cells, tissues and biochemistry rather than the individual atoms composing it. It can be thought of as a kind of trimming or pruning in which one realizes what isn't worth knowing and thereby leaves open what is. Such subtractive knowledge eliminates useless information. The rationale for this change in perspective is that assuming that all new knowledge derives from some primary, foundational knowledge risks a regress problem. Before there was knowledge there was ignorance, and knowledge is shaped out from it. Ignorance on one level can be thought of as a kind of superposition of all possibilities that have yet to be ruled out, whereas true knowledge is one substate of this overall phase space, like finding the first 1000 digits of pi in an infinite random sequence of digits.

Conveniently, because our subject matter happens to be about ignorance more than knowledge, we can more easily recruit the powerful methods and tools of information and probability theory. Thanks to strong theoretical reinforcements brought in from such fields, it is actually a more achievable task to characterize the laws of ignorance than it is the laws of knowledge. We will get to the heart of this matter in due time. For now some groundwork must first be laid out to set the stage to define the basic types of ignorance.

> Every kind of ignorance in the world all results from not realizing that our perceptions are gambles. We believe what we see and then we believe our interpretation of it, we don't even know we are making an interpretation most of the time. We think this is reality.
>
> – Robert Anton Wilson

On the most basic level like magnetic charge or any binary state, ignorance can be divided into two distinct elementary types, negative and positive.

Negative ignorance is the most intuitively understood because it is simultaneously the most ubiquitous and familiar. It is the simple lack of knowledge or information. Negative ignorance is antecedent to all knowledge in time, and is therefore a primal starting condition. Before any knowledge was arrived at there was first this kind of ignorance. Positive ignorance in contrast is false belief, it has actual mental content and is characterized by presence rather than absence.

Ignorance understood in the traditional negative sense, even without further extensions to the concept, is of substantial importance. Much can be said of it. Any time there is a problem it stands in its shadow . If you already knew the solution the problem it wouldn't be a problem. So all problem solving involves some dialectic or engagement with what is not known. Problem solving is the behavior that devises a means to resolve negative ignorance. Negative ignorance by itself, therefore, is a fundamental organizing principle for all life and decision making, to the degree that life can be described as a continuous problem solving and decisioning process that relates an organism to an incompletely perceived environment that must be explored.

Such negative ignorance can also be thought of as potential, free space, receptivity, the *possibility of knowledge*. To be ignorant pure and simple is to have ruled nothing out. To not know at all is better than to believe falsely, claims Thomas Jefferson. To believe falsely (to be positively ignorant) puts you much further away from knowing than simply not knowing or having any doxastic commitments whatsoever. And supposing one has the requisite intelligence and background understanding, information can be quickly slotted into such empty ignorance and become knowledge without resistance. Thus negative ignorance can also be thought of as potential knowledge.

Conceived in the negative, ignorance is utterly simple. Its simpleness can be deceptive however. One kernel of ignorance nestled hidden inside a plan can spell ruin to the grandest of designs. A puzzle cannot be completed with even one missing piece. In this way the relational importance of negative ignorance to existing knowledge is instructive. These connections between positive and negative factors bring forth complex interactions that are unwieldy because of all

the possible relations knowledge and ignorance can exhibit.

Knowledge and negative ignorance cannot be logical complements and reversals of each other for the fact that knowledge always has intricate representational structure while this ignorance is always purely simple and contentless. In logic, a negation is a reversal of the truth value of a proposition by applying the negation operator. The negated proposition is expressed as "it is false that p (is true)." In a state of negative ignorance, there is no such propositional attitude to be had and no proposition to be aware of to be negated. There is no reversibility of operations such that if one were to just negate the negation, reverse ignorance upon itself, you could have the knowledge it denies. So negative ignorance is more accurately the total absence of knowledge rather than its logical negation because absence is logically irreversible while negation isn't. This idea hints at the concept which will be returned to again and again, namely that ignorance and knowledge are not symmetrical inverted mirrorings of each other. Which is to say, terms like "knowledge acquisition" are misleading because it suggests that what is to be learned shares some semblance to what is known. This would imply that knowledge is out there, like an item on a table to be picked up. knowledge formation, it is to be argued is instead at its most basic level the collapse of information free parameters, in other words, it is a form of information gain in which alternatives are excluded until the fact is constrained. (More about this distinction between negation and absence of awareness from a technical perspective concerning epistemic logic is provided in Appendix I.)

So for someone to be ignorant, there is nothing for them to be aware of to negate, only a blankness. According to the principle of computational physics called Landauer's principle, any logically irreversible operation, such as erasing a bit, implies some irrecoverable information loss and increased disorder in the system's internal state or environment. What's done cannot be undone without some cost, and to negate knowledge is more like to delete it. Negation is not the same as absence, however. Lack of awareness of the truth is not an sufficient definition of ignorance either, for reasons that will be made explicit in due course.

One can also dwell in what poet John Keats called negative capability, the

power to tolerate and even draw creative sustenance from uncertainty, ambiguity, and unknowing. To dwell in such fruitful negativity is to harness vagueness and mystery, to recognize its potential, to stand in wonderment before infinity. Ignorance can be a source of inspiration, even utility; a motive to devise, speculate, brainstorm, push forward proactively in the pursuit of knowledge. Or even to wonder in undirected reverie without knowledge as the goal, to daydream, spitball, imagine without restraints. Properly recognized, negative ignorance is an epistemological motive force . It is what triggers the pursuit of knowledge, drives it toward the unknown like gravity. The constructive role of ignorance is also worth appreciating. It's not always bad. It can be a resource, a tactic, an essential ingredient to idea generation and creativity.

Positive ignorance is something else entirely. It is misunderstanding and misconception. It is false belief. To believe is either to believe that something either is or is not the case. However to believe what is not the case does not necessarily entail believing the negation of what is the case. It can also mean to believe in a parallel counterfactual world completely unrelated to any grounding in actual premises. One can believe that the Yankees will win the world series, when actually they lose it. One can also believe the Giants will win the world series, when the Yankees win it. These are distinct propositions even though believing that the Giants will win when the Yankees do logically entails believing that the Yankees will lose. However, whoever believes the Giants will win does not bother to also believe that the Yankees will lose, that possibility is excluded by the premise . One can also believe that neither team will win, but that the Red Sox will, or any number of arbitrary alternatives without there being the possibility of negating the falsehood to arrive at the truth. Believing is to be filled by a certain mental content and to be in the grips of it. To have a false belief is not only to be negatively ignorant that it is false and not the truth, but to also have an inaccurate mental representation which subjectively renders a falsified view of the world.

Though it is not directly relevant to the main argument, the notion of positive ignorance curiously resembles the concept of *avidya* in buddhist teachings (Lama Tsongkhapa, Lamrim Chenmo V3 Pp 224-267). Avidya is usually translated as

"ignorance" but it is distinct from the mere privation of knowledge that is negative ignorance. Avidya is a mental image that throws off illusions, distorting the perception of true being. In the Mahayana school of Buddhist thought, ignorance is defined as the *object of negation*, the perceptually manifest untruth that must be eliminated to arrive at the truth. This object of ignorance consists of the bare apprehension of things as individuated distinct objects separated from everything else . Ignorance on this account has an object, a *presence*. It is a confusion of the labels of an object with the object and the mistake of thinking that because language can assign distinct labels, the object that is labeled must be distinct and separate from the rest of the world. This ignorance as the illusion that things have a distinct ontological identity at all separate from the totality of being and is caused by the logic chopping and compartmentalizations of rational cognition. In other words, in Mahayana buddhism, *perception and thought itself is, or can be falsified*. The sensible world is illusion or *Maya*, falsified presence. Insight comes from distrusting or overcoming this deceptive consciousness of falsified materiality. Maya is the veil of illusion which separates oneself from everything else and trains us to think of ourselves as isolated, hungering and desirous egos detached from the rest of the universe.

Similar to the mathematical concept of the additive inverse, in which a number when added to another number equals zero, Mahayana philosophy speaks of a *non-affirming negation*, a negation which removes what is there but does not leave something in its place. The goal of the diminishment of ignorance is in other words to reach pure zero; as soon as ignorance is eliminated whatever is left must be knowledge. One must stop being wrong to have the chance to be right. To assert a non-affirming negation is not to presumptively recommend any replacement for whatever is removed. Instead, one is left with reality, what remains when all misconception and misunderstanding is shed. As roboticist Rodney Brooks said, the world is its own best model. Truth is the negation of oversaturated thought, it is the folding into and agreement with the uninterpreted, plain reality as it is.

Enlightenment in this respect is to deny that objects have inherent existence, in other words, to stop believing that one's senses and categories are the true

measure of reality and that things are what they seem. While this buddhist concept its not exactly what is meant by positive ignorance in this context, it is an interesting parallel. The claim put forward here is not so extreme as to say that objects themselves do not have distinct identities or that the senses and thought itself is intrinsically illusory, though sometimes it is true that the same machinery of representation that produces knowledge also produces illusion, it is not as though there are two separate channels that handles both separately.

For example, to tell someone that they should cut refined sugar from their diet is not to tell them what to replace it with. The implication then is that by clearing away negative false content, the result is the truth, without exactly specifying or implying what that truth is. To do otherwise is presumptuous and presupposes one is in the possession of the truth, when indeed it is entirely possible to simply swap in another false idea in the old false idea's place. The desire to understand an unknown truth which itself is motivated by the desire to remove ignorance, implies that one can recognize the truth when it is perceived which is a contradiction. Otherwise you already have what you sought. How else could you recognize a truth that is unknown? This is a kind of eliminative cognition proceeds by vacating falsehoods and extraneous details until only the truth remains in balance with the world. For these reasons, it can be argued that knowledge acquisition is more accurately ignorance removal, and this are no mere tautological reversal but an entirely distinct statement.

Appearances can be misleading. *In The Case Against Reality: How Evolution Hid the Truth From our Eyes* cognitive scientist Donald Hoffman pitches what amounts to a scientific variant of Mahayana philosophy. He compares conscious perception to a kind of virtual reality interface, which models not how things are in themselves but rather emphasizes the salient aspects of which are useful to the organism, or which fabricates purely subjective overlays that override and manipulate objective states in service of biologically centered interests. Both the Mahayana and scientific arguments portray conscious perception as a kind of positive ignorance, an internal representation which projects a contorted picture of the perceived.

One might respond that positive ignorance is always modulated by not knowing—a person experiencing an illusion who knows it's an illusion is not under its sway. It's true that positive ignorance is always regulated by negative ignorance, and it is cancelable by destroying the negative ignorance which regulates commitment to the false belief. However what is missing from this critique is that illusions have the *power to deceive.* Illusions, propaganda, misinformation, can induce ignorance that was otherwise not there. It's only by thinking through it that its capacity to induce ignorance is disarmed. Negative and positive ignorance can be combined in the sense that what one doesn't know and what one falsely knows can scaffold and reinforce each other. This is what I call complex ignorance.

"A great deal of intelligence can be invested in ignorance when the need for illusion is deep." -Saul Bellow

The common sense understanding of ignorance is what can be termed simple ignorance. Like negative ignorance, it is the simple lack of knowledge. Not knowing how to perform calculus or how to tie a bowline knot or not knowing what the capital of Mongolia is simple ignorance. All negative ignorance is simple ignorance. More essentially, simple ignorance is bounded, it is a discrete, atomic component. Simple ignorance can also be a positive misconception, the belief that the capital of California is Los Angeles, for example. The key point is that unless such simple particles of ignorance are incorporated into a larger belief system, they are atomic.

There is also what has been called misbeliefs (Dennett and Mckay, 2009), and belief systems which can thought of as *complex ignorance*, ignorance that is built out and articulated from many such simple components. It is ignorance that is systematic, consisting of relations between positive and negative ignorances. This is not that someone doesn't know something, but that they are comprehensively wrong. Complex ignorance can be likened to a combination of simple ignorances. Moreover, actual knowledge can interact with such complexes. It is possible to have a mixed state in which both positive and negative ignorance and

knowledge are involved in what is essentially an ignorant formulation. In essence, complex ignorance is marked by its *representational content*, it is systematic and joins together multiple individual atoms of ignorance in a schema.

Negative ignorance often plays a regulating role in maintaining misrepresentations in that an otherwise rational agent if made aware of it would then revise their thinking or otherwise reject it entirely. Positive simple ignorance, however is what makes up the content of the misrepresentation alongside misapplied knowledge. As will be seen later, negative ignorance can structure the available options for thought, lending the appearance of false wholeness to some ideas. If one is subject to second order ignorance and does not know one doesn't know something, the what one knows one knows and knows what one doesn't know seems like the whole landscape. Running with this landscape metaphor it is not only that the map isn't the territory, but the positively ignorant misrepresentation is like an antiquated map that depicts the borders of countries that are no longer up to date or a flawed map from the 16th century that leaves out the Americas. One can still navigate according to such a map just as good as an accurate one but it will not lead you to the right places.

One might wonder whether there is a calculus or rubric for evaluating what differentiates merely incomplete, though valid knowledge from complex ignorance. In real world scenarios knowledge always to some extent faces down ignorance. There is nearly always some influence of ignorance on an agent's cognitive state. Yet in good conditions this ignorance is an information deficit, an imbalance between the internal information state of the agent and the external information state of the environment. . As such there is a defined gradient. In principle, although not always in practice, this imbalance can be corrected. The likelihood to do so is much lower when thinking is guided by complex ignorance. Complex ignorance occurs when there is less of a clear demarcation between internal and external information and the agent incorporates mistakes or misinformation into their mental model. True knowledge, one might say, is linearly separated from ignorance, it is knowledge through and through and ignorance is on the other side of the separator. It can also occur when the

assumptions driving the model itself are grounded in a false premise. Just as it is possible to draw any conclusions from a valid but unsound argument, so is it possible for ignorance to steer thinking. Similarly, some premises of a valid but unsound argument be true, while others false. The requirements of a sound argument is that *all* premises are true as well as the conclusion. On the basis of classical logic, therefore, it can be seen that much complex ignorance consists of such valid but unsound arguments. They are valid enough to be mistaken for knowledge, and some of the terms might even be true, but it its nevertheless inaccurate.

Complex ignorance takes many forms. It is ignorance that, somewhat paradoxically, often has intelligence invested in it. An explanatory theory which is ultimately false is one example. Darwin developed a theory of genetic inheritance called pan-genesis which for all its ingenuity was misconceived. The key point is that it was a theory, it had definite representational content and even clever reasoning behind it. Or consider warfare. War is arguably an irrational thing *done* intelligently and rationally. It's been said jokingly that military intelligence is a contradiction in terms, but this quip like so many jokes hides a subtle truth. War is a form of solutional ignorance, politics by other means as von Clausewitz described it, implying that better, peaceful solutions were untenable. Tremendous amounts of intelligence are contributed to the organization, training, armament, strategy and tactics of armed forces yet war amplifies human suffering and destruction to the utmost degree.

Perhaps few examples embody complex ignorance as good as conspiracy theories. What is a conspiracy theory? It is, above all, a theory—it has a framework that asserts a set of propositions and a set of relations between those propositions. More than simply products of stupidity, conspiracy theories can even be clever, ingenious, inventive, but still fallacious. While it is not self-evidently true that conspiracy theories are necessarily false, they often are wildly so. A person who embraces a conspiracy theory is not considered ignorant because they do not know or deny the truth, although that is a necessary condition for belief in one. They are ignorant because they accept a complex web of misrepresentations as the truth. They picture to themselves a false world.

Their ideological projection overrides empirical evidence or explains it a way as deceptions in league with the conspiracy. They are ignorant in both what they choose to include and exclude from their story. Such theories are arrived at through a line of misguided reasoning and delusive pattern recognition. They provide a payoff for the conspiracist, sating their otherwise intolerable uncertainty. Given the argument that thinking is all about uncertainty reduction, it goes to show that uncertainty of the subjective sort can be equally abated by both false and true explanations. Such a theory is not an absence of truth or even a chaos of randomness, it is a false order. Belief in a conspiracy theory isn't necessarily willful ignorance, because willingness to believe falsely implies that the conspiracist is intentionally forcing out reality. They believe the conspiracy is knowledge and they are able to construct a seemingly consistent system from thin air to justify it. DeNicola calls such complex ignorance false knowledge, although I would argue against this term because knowledge must always be true. One suspects "false knowledge" is used in place of what I call complex ignorance because false knowledge is a contradiction in terms as all knowledge must be true. It is assumed that ignorance lacks representational content while knowledge clearly possess it. This not strictly the case.

Distinguishing between genuine knowledge representations and complex ignorance can be a great challenge. The history of ideas is littered with discredited theories that were once embraced as true. Aristotle advocated a theory of spontaneous generation in which life manifested from inorganic matter constantly. For quite some time one of the popular theories in the physics of light was the concept of the luminiferous aether, the supposed medium in which light propagates. The earth was once believed (and sadly among some still believed) to be flat. None of these theories were simply the absence of knowledge, they were products of thinking. And in their time it was considered the reserve of the learned to believe and study them. It goes to show then how the line between knowledge and ignorance is blurred, how ignorance can wear the outfit of knowledge. Surely much of what is believed to be knowledge now will eventually face the same fate as these superseded explanations. To appreciate how knowing happens in real world settings, the wider picture and note how

ignorance and knowledge play out on the same field and are bound by the same rules. knowledge as it exists in real life, as opposed to idealized conception, is grainy, rough edged, and imperfect rather than some platonic pure form.

A major rationale for this notion complex ignorance is that it provides a fresh way to study misrepresentation. If ignorance can have content, and can be fit into representations the same way that knowledge is, then this blurs the line between knowledge and ignorance. They are more continuous than may seem at first glance. Imperfect knowledge is not exclusively incomplete knowledge, it can also be knowledge that is concatenated with ignorance. It can also be knowledge that is applied ignorantly, or even ignorance that is backed up by knowledge. The idea that ignorance can be combined into structures at first chafes against common sense given that ignorance is typically understood as nothing. If beliefs can be ignorant, and beliefs can be tied together in networks, then there is nothing strange about this commitment. However ignorance is not merely what its left out of imperfect knowledge, it involves the connections knowledge that knowledge has with contents that are not knowledge. Otherwise there would be no way to distinguish someone who is *un*informed from someone who is *mis*informed. The misinformed is more far gone than the merely uninformed. The uninformed person simply lacks the correct information and they still have the opportunity to gain it, they are available to take on new information. While the misinformed constructs a representation of a distorted reality. It would be better for them to be uninformed. This would actually bring them closer to knowledge. That ignorance can dress up as knowledge and directly impinge on it proves to be an important fact when appreciating how knowledge is obtained and operates in the gritty, confusing real world.

Descartes coined the term *res cogitans* , thinking matter, to describe the substance that thinks. This was the irreducible kernel of certainty left on the bottom of the pan when everything else was boiled away by his method of doubt and affirmed by the dictum cogito ergo sum. The “I am” that thinks is this thinking stuff. In one stroke, Descartes pinpointed the attribute which would be the sun at the center of conventional epistemology’s affirmation-centric cosmos,

around which orbited the manifest and positive arguments about knowledge it asserts. These positive ideas of knowledge originated from the affirmative nature of the reasoning faculty, which actively posits and asserts what is the case .

I would like to suggest a spin on Descartes' turn of phrase. There is something like the stuff of ignorance, or to be grammatically precise, the unknowing thing: *res ignorans*. This is not the same as regular non-thinking stuff, plain, inanimate, inorganic material, sticks and stones and chairs, although it may relate to them as its object. Neither is it exactly the reverse side of the thinking thing, the thoughtless mind, although the thoughtlessness does play a role in it. So it is not a simple inversion of res cogitans. Rather, it is the shadow and the backdrop, the context informing thinking matter, the space surrounding the sun. It is everything Descartes doubted in order to arrive at the thinking thing. It is also tucked away within this thinking thing, a secret stowaway that his method of doubt overlooked. It's ignorance as a substance.

That ignorance as a thing is at least as real as thinking as a thing and exists beyond the mind might be seen as bizarre. How is it mind-independent? Negative ignorance is not like thoughts, as it is not present in the mind as a content, so it should not exist at all. Negative ignorance is *intentionally inexistent*. It sounds absurd. Ignorance is at most abstract, and what makes an entity abstract is that it lacks substance.

It may take some work to convince the reader that ignorance has its own independent reality. That is the current goal. Common sense dictates that ignorance exists only in relation to knowledge, and by extension the mind. And what existence it has is faded, counterfeit; because it represents either belief in falsehoods, or merely the absence of knowledge. knowledge certainly can't exist without a subject to know it. Almost tautologically, this would mean that ignorance must be mind-bound, there must be unknower, to not know. What gives?

These concerns will be addressed in turn, but first, I'd like to preface what follows by saying that this way of thinking about ignorance unlocks a useful philosophical toolkit for reasoning about this important topic. All the techniques used to reason about actual objects, or at least epistemic objects can then

transitively apply to it. It is like recognizing that what we perceive as empty space to the naked eye is actually replete with force fields and cosmic rays. There is something to the darkness.

There are striking similarities between this externalist account of ignorance and C.S Peirce's doctrine of *objective idealism*. For Peirce, there is no true distinction between mental and physical states. The mind and its contents are directly continuous with the material world and mental things such are ideas are physical things and physical things mental things. This continuity also means that mental states such as doubt are reflected in the state of the world at large. Peirce is committed to a fallibilist position, in which, "knowledge is never absolute but always swims…in a continuum of uncertainty and indeterminacy." (Peirce, 1879, p356). Such wording ("swimming in") is striking because it suggests that the uncertainty in which the mind deals is an objective property. The universe, Peirce contends, is inherently indeterministic in nature, though given time it evolves into a more systematically arranged order . Like a funnel, the initial primordial uncertainty is gradually channeled into "habits", regularities that tend to stick to a baseline of most probable behaviors. This tendency for things to drift toward habit is true of both mind and matter. The laws of nature therefore "fall into place" from an initially disordered chaos, rather than having been there at the start.[1]

Ignorance is certainly tied to mindedness, and that is where it has meaning. If tychism has some substantial basis in reality, and quantum mechanics might suggest it does there are good reasons to believe that its existence is not strictly dependent on minds. For example, consider the fact that when we are ignorant of something, it is not part of our mind, yet we still speak of a known unknown or the unknown unknown as a fact. What are we speaking of in this case, if not something that is not of our minds? One answer is that it is an example of intentional inexistence, like the thought of a unicorn which is not tied to any real world object. Except that's not strictly the case with ignorance. It is tied to facts

[1] Later we will see how recent research into what researchers are calling the autodidactic universe follows in Peirce's footsteps in viewing the laws of nature as gradually emerging from an original inchoate state.

about the real world and the way the universe at any given state of time is constituted. We are capable of correcting ignorance, of responding to it by learning what hides behind it, as if by lifting a veil. As Rescher (2009) argues, ignorance of things that do not exist or that are not facts is meaningless.[2] It is an idle point to assert that it may or may not be a fact that if Archduke Ferdinand wasn't assassinated WWI wouldn't have happened; there is no grounds for verifying this claim. Yet there is also this sense in which a shell of possibility informs all thinking when uncertainty is involved, and that when plausibility and possibility are added together you get probability. Through this transformation of the possible into the probable knowledge becomes more definite as it emerges from out of a haze of uncertainty.

Through learning the object that is veiled by a curtain of ignorance is discovered and transitions from unknown to known at various increments. Yet ignorance in this case is not a mental state. And while ignorance is associated with the relation, or rather lack of relation, between the mind and the hidden fact, we are ignorant of it for the ultimate reason of the causal history of the entire world, the way that things shook out in the entire history of the universe, including especially that course of events that never put us and that object in range of discovery. In this sense the universe itself, the sum of all causes and effects, impressed ignorance on us, it is an effect of how the world is in relation to the agent. For these reasons it is conceptually useful to think of an unknown as a sort of buried object, as a thing tucked behind a veil of uncertainty, rather than a nothingness that is replaced by knowledge .

More to the point, ignorance is bound to the limits of information availability and accessibility which is largely determined by immutable physical laws. We cannot detect information that exists within the depths of a black hole, at least not directly. This is not due to the limitations of our minds but the limitations of

---

[2] Rescher also discounts unknowable counterfactuals as valid forms of the unknown, because the factuality of these claims cannot be asserted. Here I disagree, and prefer a modal realist view of possibility as inherent to the nature of ignorance. The more we know, the more we exclude excess alternative hypotheses

physics as it constrains information flow. This fact about what information is like a shadow with light streaming past it. The shadow is still constitutive of the visual scene: its reality cannot be written off. If Kant is to be believed, we cannot know the Thing-in-Itself, the inherent essence of external reality independent of the mind's categories of reason and perception. This inability to know the mind-independent essence of the world, then, is itself an *objective fact*. This unknowable essence is, indeed, the ultimate fact. Which is another way of saying that objective forces which structure the facts of the world permit us to know what we can know and rule out what we can't. Ignorance happens to a person, just as how getting lost happens to them.

Even more to the point, if ignorance is inseparable from mental relations, one wonders how a newborn passing into life or a dying person passing into death relates to their mental oblivion before life or after death, when there is no mind to do the relating. Aren't they ignorant of the moment before their birth or the moment after death even though their mind does not exist? Is this mindless obliviousness any different from the insensate unconsciousness of other kinds of unawareness such as a catatonic coma or dreamless sleep except that it is taken to the extreme? If so, then ignorance would have been present even though they were never born or are no longer alive. Which is to say, a possible knower in this world is denied the possibility of knowledge, a situation which defaults to ignorance. If there is no possibility to know, then it is only possible to be ignorant. More precisely, since there is no "to be" it is only possible to be ignorance, to totally negate the conditions of possibility for knowledge.

A hypothetical universe which lacks any conscious, sentient minds within it lacks both the existence of knowledge and *subjective* ignorance. Only objective ignorance, as the absence of the conditions of possibility for knowledge, remains.

One can write:

$\exists s \in U, (Ks\beta \oplus \neg Ks\beta)$:

there exists at least one subject s in the universe U such that s knows or does not know β. This describes a world such as ours that contains at least one knowing or ignorant subject and contains the possibility of knowledge.

Uncontroversially, in this world it can be stated

$\Diamond\neg Ks\beta \rightarrow \Diamond\neg K(\beta)$:

it is possible s does not know β therefore it is possible that β is unknown. Then you can also say that in such a universe:

$\Diamond Ks(Ks\beta \vee \neg Ks\beta)$:

it is possible that s knows it knows or knows it does not know β.

One can also say of a universe without at least one subject:

$\nexists s \in U, (Ks\beta \oplus \neg Ks\beta) \rightarrow \neg\Diamond Ks(Ks\beta \vee \neg Ks\beta)$:

This reads: there is a universe where there does not exist at least one subject such that s knows or does not know β. Therefore it is not possible for at least one subject to know it knows or does not know β. But this lands you with:

$\exists U, \Box \neg\Diamond K(\beta) \rightarrow \neg K(\beta)$:

there exists at least one universe where it is necessarily impossible that β is known. Therefore it is not known that β.

This clearly demonstrates that even in a universe without subjects of knowledge, there can still be facts of that world such as β that are not known. Ignorance precedes the possibility of knowledge. Another way to put this modally which I won't belabor , is that if there is a universe where there is no knower to know β, then β is unknowable.

This formula expresses fundamental, mind-independent ignorance. The removal of the possibility for knowledge defaults to ignorance. Ignorance is prior to knowledge. If it is impossible to have knowledge in a mindless universe because no knower exists to know it, then ignorance of the facts of that universe must still exist, because there might at least be a possible knower that could know it. But there are no possible knowers available to know of the facts of that universe, thus any knower is ignorant of it. But anything that exists has the possibility of becoming known, if, for example, within such a universe knowers eventually evolved and find it out. Certainly every fact that existed before the first knower was born was unknown, or to be exact, unknowable at up until that point of time at which it could be known.

Our subjective ignorance is therefore constrained and permeated by objective factors which induce it. Objective ignorance can be thought of as a sort

membrane or effluence that emits from these, ultimately physical, causes. From a personal standpoint there is a real sense in which ignorance is an imposing outside presence, a fact in itself rather than a fact about the limitations of our minds. We find it natural to think of it as an effectual, active, meaningful thing. Claiming that ignorance is in some sense objective lends credence to this ignorance-phobia, otherwise we would just be afraid of our own minds, or else afraid of our intellectual flaws. We aren't afraid that our knowledge is inadequate, at least not exclusively. We are anxious of ignorance and troubled by uncertainty because the way the universe is condemns us to a fate of ignorance always to some degree no matter our intelligence. Ignorance is the unseen tiger in the bush. It's the fact that the world seems destined to throw more at us than we can anticipate or understand which is the true source of such epistemic anxiety.

Even if ignorance is of course not really "out there" floating around in the in ostensively definable blobs, it is intellectually and pragmatically compelling to think of it as existing beyond the scope of the mind's limitation. It has an independent logical reality as something embedded in the structure of the world. Thinking of ignorance as inextricably attached to a subject leads to inconsistencies. It preempts more robust methods of analysis. In terms of net gain in explanatory power and conceptual clarity, much is to be had by viewing ignorance as substantial at the mere cost of a little common sense.

It's helpful to point out too that there is a deep literature and debate over the reality of abstractions (see Zalta, 1983). It goes all the way back to Plato, who championed the idea of the objective existence of abstractions and more or less brought the idea to the absolute. For Plato, the Forms are *more real* than sensorily perceptible objects. So is by no stretch an unheard of suggestion that abstractions may possess independent, or even fundamental, existence. Ignorance is clearly an abstraction in that it is intangible. While the supposition that abstractions have separate existence is not necessarily true, neither is it necessarily false. Let us also not overlook the fact that we now live in a world dominated by abstractions of a computational sort, database schemas, network protocols, user interfaces, programming languages, cryptocurrencies: the indisputable reality of

intangible software reigns supreme. And it almost goes without saying that are our own states of consciousness through which all experience is mediated is intangible and lacks properties attributable to matter (extension, mass, etc).

What then is the difference between the substantial and the abstract? In the usage stressed here, I mean effectual—able to influence matter or cause behavior. Plato's abstractions are substantial in that they inform all perceptual instances. In the same way, ignorance is substantial in that it structures what is knowable and conditions what is known.

It's widely understood that quantum physics describes its phenomena as indeterministic and probabilistic. While this attribution of native uncertainty to quantum objects has its skeptics (most famously Einstein), it says something profound that we appear to have no other way to conceptualize these things without resorting to probability. Is our need to use probability reflective of the limitations of our minds? Are there truths that can only be half-grasped in this way? Is the quantum physics of today just a half-measure, an incomplete part of a perfectly deterministic formulation that nobody is smart enough to figure out? Or is there a deep structural reason embedded into the fabric of physical reality that forcibly hoists this way of thinking on us? Fittingly, the question itself feels unanswerable. Here lurks the shadow of objective ignorance; it is the abstract correlate to the indeterminate physical fuzz of quantum uncertainty.

There's also the temptation to parallelize knowledge and ignorance, to view them as night and day, two sides of the coin, evenly proportioned reversals of each other. This is not strictly true as they are demonstrably asymmetric. Ignorance is conceivable as potential or possible knowledge, but this is to assume that the knowledge is already latent, either lying out in the world "to be acquired" or somehow already in the mind waiting to be activated (such as Plato's "unforgetting"). And it can't be ruled in that everything that is unknown is knowable, this fact itself is arguably unknowable, which alone specifies an asymmetry. Some ignorance in convertible into knowledge, but there are almost assuredly unknown unknowns which are not potential knowledge because they are not discoverable.

knowledge and ignorance are asymmetric for other reasons such as in terms

of quantity. knowledge is precious, extraordinary, rare. It takes effort and energy expenditure to acquire it whereas ignorance is freely given. Ignorance is the most commonplace and mundane thing in the world. knowledge is a low probability, exceptional state, while ignorance is profuse and endemic.[3] It would therefore make sense in purely quantitative terms that knowledge would be pooled in the reservoirs of our minds while ignorance roams freely, entangled with the fluctuating entropic volatility and chaos of an exceedingly complex and perhaps intrinsically indeterministic world. knowledge it can confidently be stated *does not* have mind-independent reality.

Another problem is how to accommodate this objective perspective for more rarified positive forms of ignorance, such as willful ignorance, misbeliefs and the like. These clearly have a mental predicate (willful, belief-ful). This problem is also explained well by Peirce's objective idealism, which views the same operative principles at work throughout both mind and matter. Minds are not magically sealed off from the physical world. Everything subjective is objective in the strict sense that whether it is a mental state, a true or false belief, an arbitrary opinion, or a delusion, it is still a fact occurring in the world just the same. Minds are, like lightening conductors, focal points, contact sites for everything to do with with knowledge and ignorance.

This story about the objectivity of ignorance is a hard pill to swallow. However, just as David Lewis's modal realism stretches far beyond the limits of common sense, he nevertheless made a good case that there are valuable insights to be gained by viewing possible worlds as actual and that the benefits might even outweigh the costs. It opens a path to a philosopher's paradise, as he called it. One might fear that inflating ignorance to such a monstrous scale creates a philosopher's hell. But just as Lewis does with possibility I argue that, although

---

[3] There is some debate whether scientific knowledge is approaching its maximum and that most or all of the major scientific theories of the universe have been discovered. This is problematic for two reasons. This assertion is on its face unprovable, because it cannot be demonstrated how much is unknown in the same way that a fish in a pond cannot know what is above water. Second, a distinction can be made between general facts and specific facts. Ignorance of specific facts can still be vast even if knowledge of general facts is thorough.

counterintuitive, accepting ignorance as substantial, as a fact of nature rather than a merely notional concept, pays dividends in terms of philosophical utility.

How does it? The short answer is that it allows us to speak and think of the set of all Ignorance and each instance of ignorance as first class logical objects and logical types with terms and quantifiers and all that good stuff, for which an organized method of reasoning such as a propositional calculus might be developed. Essentially, it helps us map out the territory of the undiscovered and reason about its relationships, a useful thing indeed. It provides a whole new way of looking at probability theory, such that probability measures actual quantities of ignorance rather than a merely notional expectation value. It paves the way towards the creation of an extended epistemic logic that can apply to ignorance-objects above and beyond mere negation statements. It's even conceivable that mathematical reasoning might be used to study it, in at least a primitive sense of applying arithmetical operations such as adding or subtracting ignorance, or inequality relations such as less than or greater than, although it is unclear how definite quantities could be assigned to unknowns. ("You have 1300 ignorands worth of ignorance on this topic!") So it makes it more immediately obvious to restrict it to logic. When all is said and done, how can being more precise in thinking about ignorance not be a good thing?

Of course, the proof is in the pudding and what is being sold here has not yet been made. Still, in what follows the outlines of this neg-epistemic project will take shape.

What really drives the argument home comes later, when it is shown how ignorance and entropy are deeply connected. So connected, in fact, that if it is accepted that entropy has objective existence then so must ignorance. The reader who wishes to learn more about how I counter doubts about this argument is invited to jump to counterpoint #2 in the Appendix I.

Is ignorance continuous or discrete? At a glance it's not exactly clear what this means. However if ignorance is substantial, it makes sense to think of it in quantitative terms. We can speak of ignorance in two senses, what one might call global or local, or better, ultimate versus proximate. Ultimate ignorance is essentially the universal totality of all ignorance, everything that is unknown,

"the unknown" itself, as when we speak of outer space as "the great unknown" rather than a collection of isolated unknown facts. This is opposed to a specific local, proximate unknown, "an unknown." Essentially ignorance can designate an indefinite or definite object, a category or a member of that category. Ignorance can mean a certain blindness to knowledge in general or to a specific case. When Plato says the ignorance is the root of all evil, he does not mean not knowing how to play the violin is the root of evil, it is the larger phenomenon of Ignorance that is. For clarity I will refer to this ultimate Ignorance with a capital I and proximate in lower case.

One might naturally conceive of the ultimate-proximate distinction of Ignorance as a set of sets relationship. The question then becomes whether or not there is a terminal subset within the ultimate superset, containing one atomic element, the tiniest conceivable speck of ignorance, whatever that means. It might be that Ignorance is a flat set, containing only one level. Not knowing how to play the violin, or fix a bike, or chemistry might all be definite, whole, and indivisible facts. However looking at it carefully those ignorances may imply smaller, component ignorances, not knowing chemistry means not knowing many separate facts of chemistry, etc. Objects are structured into part-whole relationships, and one can be ignorant of the whole or some combination of the parts. Supposing the universe itself is finite, we can be sure there is an upper bound to Ignorance, that there must be a universal set of *all* Ignorance, whether it's the set of things that no one knows or a person's specific "ignorance fingerprint". Though whether this set contains finite or infinite first order elements may depend on whether the universe is finite or infinite, as presumably an infinite universe would entail an infinite number of unknown contingent facts about it and vice versa. If the highest order subset terminates in a set which contains elements which are empty sets or not sets, then ignorance has a fundamentally discrete, atomic quality. In other words, it would then be the case that it is theoretically possible to know everything about a given object.

The case can be made that knowledge has such a discrete floor, after all knowledge is indisputably finite, and knowledge is structured in a way that is far more salient than ignorance. But even knowledge of mathematical infinity, like

knowledge of Ignorance, does not mean a total encompassing knowledge of it. It's more like a synopsis. Given that knowledge and ignorance are not symmetrical homologues, we can't infer that it must have the same logical form as knowledge.

I'd argue that it is impossible to establish grounds for the discreteness of Ignorance despite the utility in thinking of it as such. Except this has more to do with its relationship between knowledge than it does the mathematical properties of continua or discrete quantities. Discreteness implies a level of granular knowledge of the specified "level" of Ignorance which immediately destroys it and replaces it with knowledge. As soon as we reach the bottom of ignorance and find out there is no more we don't know, we have parsed it and therefore acquired discrete knowledge. "Ignorance," wrote Oscar Wilde, "is like a delicate exotic fruit; touch it and the bloom is gone." Until knowledge replaces it, Ignorance is therefore bottomless.

To take a concrete example. Assume we are talking about one person's ignorance fingerprint for clarity— the totality of what they don't know that you get when you subtract what they do know. Their overall ignorance contains ignorance about playing the violin. But this may entail not knowing how to control fine motor skills in a precise way, how to learn how to play the violin, how to learn how to learn how to play the violin, how many ways to play the violin, when to learn how to how to learn to play the violin, where to learn how to play the violin, how those who know how to play the violin know how to play the violin, how to play all the compositions that a violinist can play, what it takes to be a master violinist, how to perform playing the violin in a symphony, and any number of interminable details. These skills in turn suggest a set of sub-tasks about which one is ignorant. Furthermore not knowing how to play the violin is nothing like not knowing how to prepare a soufflé, though at the same time there is a clear featureless uniformity across all ignorance. They are ignorant of both skills for different reasons. From the perspective of the unknower, there is no bottom to their ignorance, it is a featureless, a pure blankness. Each of these "subordinate" ignorances about subtasks tied to the "overall" ignorance of not knowing how to play the violin may be sub-dividable to a final limit, but once

we have identified that bottommost level we have begun to *know* something about the subject matter. That is, we learn how much there is to know about it, at which point, the unknown is replaced by the known: knowledge is discrete but ignorance continuous. Because the ignorant subject is by definition unable to specify the grounds by which they are ignorant without thereby decreasing it through gain in knowledge, their ignorance must be taken as continuous in scope and homogeneous in consistency.

Leibniz proposed what he called the Law of Continuity in the course of developing the infinitesimal calculus. He summed up the principle with the aphorism, natura non facit saltus, "nature makes no jumps." The idea is that all natural phenomena occur within a uniform continuum, that there are no "interruptions" in the workings of nature. Summing up the rule, Leibniz said,

> In any supposed continuous transition, ending in any terminus, it is permissible to institute a general reasoning, in which the final terminus may also be included.

That is to say, in any continuous process, any supposed end point itself counts as a continuous process, and so on for any of its points. From this idea Leibniz derived his notion of an infinitesimal number, a number that is infinitely small but greater than zero. In a similar vein, C.S Peirce thought that the empirical existence of continuous quantities meant that nothing could be known with perfect certainty. After all, he reasons, the value of a continuous quantity can only ever be approximated, the derivative of an infinitesimal is itself infinitesimal. Consistent with this notion of continuity is the idea that there is no sharp division between mental and physical things. Indeed, Peirce insists, this continuity demands that the mind cannot operate on anything besides ideas, because for it to be otherwise there would have to be a discontinuous jump between the types of entities the mind works with yet this cannot be the case if all mental contents flow together continuously. Thus, "the reaction between mind and matter would be of no essentially different kind from the action between parts of the mind that are in continuous union, and would thus come directly under the seat law of mental association…" (p.353). Accordingly, the law of mind

applies equally to matter. Through this means, mental states blur into physical states, which opens the door toward there being a kind of uniformity between properties usually attributed to mental states, such as uncertainty, and external states of matter. Such a line of thought leads to such profound implications. The most relevant for the present purposes is that it makes this assertion of objective uncertainty more plausible. Indeed, Peirce's idiosyncratic theory of tychism is a clear precursor and original statement of the concept.

If Perice's view on the intrinsic indeterminacy of the universe is to be taken seriously, and if he is right that the mind is continuous with physical states—that ideas are essentially of the same type as physical objects despite their abstractness—then the claim that ignorance goes beyond the mind and is intertwined with nature's workings on a fundamental level starts to add up. The internalist-externalist distinction is arbitrary because there are no disruptions in a continuum. There is always more to know, and there is always more ignorance, because the chance-ruled, bottomless nature of the world makes its possibilities inexhaustible, and these two facts derive from one and the same truth. All thinking, according to Peirce, is an act of approximation. The discovery of truth approaches the asymptote of a determinate truth value, but there is always something a bit more to it that seemingly recedes in proportion to which progress towards its limit gains. The order that is perceived by knowledge, says Peirce, *is* the order of knowledge. This order consists of "congealed" substates of this broader, aleatory configuration space, resting states that are reached as the universe evolves from a more chaotic to a more organized form and acquires "habits." What is empirically knowable, those phenomena which appear to be the byproduct of deterministic natural laws, are actually the exceptions, contingencies at odds with the prevailing rule of chance.[4]

It's arguable whether continuity holds true of the physical world as modern science conceives of it today. Quantum physics interprets the nature of the microscopic to be fundamentally discretized and composed of units, or quanta,

[4] As to be expected Peirce's views on probability align with the objective, frequentist school of thought that views it as an actual property of the relative frequencies of events. As opposed to the subjectivist, Bayesian school that interprets probability as the assignment of variable degrees of credence to beliefs.

of finite energy that can be assigned a complex number representing a probability amplitude. Theories of quantum gravity even stipulate that space has a smallest unit and that gravity is quantized (Smolin, 2002). At the same time strangely enough all these discrete particles are simultaneously continuous waveforms, hence wave-particle duality. Furthermore all physical phenomena are thought to be "excitations" of fundamental fields of electromagnetism and so on. But that is somewhat besides the point. Leibniz also developed the theory of monadism, which states that nature at the fundamental level is composed of indivisible entities called monads that express a perfect simplicity. The world, Leibniz wrote in his treatise on the concept, *The Monadology*, was such that "each portion of matter can be conceived as like a garden full of plants, or like a pond full of fish. But each branch of a plant, each organ of an animal, each drop of its bodily fluids is also a similar garden or a similar pond." And so on, turtles all the way down. This infinitely divisible fractal continuum could be measured in terms of infinitesimal quantities, and so it was not beyond the grasp of reason.

Long since abandoned in favor of empirical physics, monadism has gathered dust in the forgotten annals of the history of ideas. Though perhaps it can be resurrected for the study of ignorance.

Ignorance, conceived in its universal totality as all that is unknown, is essentially monadic. It is an perfectly uniform continuum, ignorance all the way down, and everywhere is still ignorance, in a perfect, utterly simple and undifferentiated form. Until, at least, it stops being ignorance. The limit of ignorance is knowledge. This is because its bottom, supposing it is unknown, is necessarily indefinite and it has no fine grained low level foundational structure. It is infinitesimal. Once a limit to Ignorance is specified it is replaced with knowledge at whatever point the discovery was made. Proof of this idea can be demonstrated by the reduction to absurdity of the opposite claim: Ignorance is decomposable into discrete atomic elements. At first this claim that ignorance is not discretized might be seen to contradict the earlier claim that simple ignorance is elementary. However, it is only meaningful to speak of elements of ignorance, the one thing that one doesn't know, if and only if there is some relating knowledge to index it. The only discrete ignorance is known unknowns, and it is

rather the knowledge that identifies the unknown, not the unknown itself, which has definite discreteness.

Suppose that ignorance, taken again as the sum of all that is unknown, decomposes into discrete units. Such units must have a definite value. This is intuitive, we say "I didn't know *that*" and the use of the indexical pronoun "that" implies that it isn't the totality of all Ignorance referenced but a discrete bounded portion, a unit. Granted, whatever this unit of ignorance is, it isn't quite a unit, because by virtue of being unknown its extent is decidedly vague. Yet how exactly would one define the smallest possible indivisible unit of ignorance? Perhaps one is tempted to use the same measurements as information, the bit. The object of one's ignorance is ultimately a set of bits the values of which have not been determined. But a bit is an answer to a question, it is not the question. It represents a terminal decision state. A bit is part of a signal which communicates that a indeterminate binary state has flipped to a definite value. It is this act of flipping which constitutes the informativeness of the signal. We can't measure Ignorance in bits, because to do so is to instantly swap proximate ignorance for knowledge by measuring the information state of the previously unknown fact. To measure is to know, so ignorance must be immeasurable because to measure is to eliminate it.

Perhaps one might go a step further and claim the measure of ignorance is a qubit, a quantum bit, which is analogous to a bit except rather than occupying one of two definite binary states it occupies the superposition, both of them, at once. A qubit has a probabilistic multivalent character that is attractive as a model of a unit of ignorance because of its indeterminacy prior to measurement or use in computation. The kicker here is that quantum information while in superposition *itself is continuou*s, because quantum states are governed by continuous wave functions.

According to the wave function principle measuring the quantum state causes it to collapse into a single state, to *decohere*. (Measuring here can simply be taken to mean the influence of one physical system by another physical system.) Simply put then, to measure the qubit, to *know it*, is to destroy its normal multivalent state which is indeterminate and spread out, and in which

uncertainty prevails. That is to say once you measure its value it stops behaving in the way it typically does which is in terms of probability amplitudes. But again, a common definition of probability is the *quantification of uncertainty*, and as we will later see in much detail, uncertainty is an important type of simple ignorance. Presiding unmeasured in a state of ignorance, the quantum bit's states are described by continuous functions that cannot be labeled with a discrete value until they are measured. To know, therefore, *is to make definite*. "We kill to dissect," a poet once said. In similar fashion we destroy ignorance to know. Ignorance is instantaneously swapped for knowledge as soon as we scope out the extent of what is unknown. Except the ignorance that has been replaced by knowledge did not share this property of discreteness because in a state of unknown unknownness its limits are necessarily indefinable. In a state of absolute ignorance, it is impossible to distinguish pieces of ignorance from the whole. The proximate unknown was in other words indiscernible from the ultimate unknown, which itself is immeasurable and unbounded. Ignorance, unlike knowledge, has no joints.

The gist is that if ignorance were divisible into terminal discrete units we have no means of representing those units as definite quantities. Even if we were to rely on qubits, when they are in superposition they are described by continuous quantities, as all states smudged together. Though quantum mechanics describes its objects in terms of probability amplitudes, it would be a mistake to say that the amplitude is identical to the quantum state being described. Quantum objects are not "made up of" probabilities. There is however this probabilistic "membrane" which seems to superimpose over quantum states, and the question is whether this is a projection of the mind and its only means to conceive of what's going on, or if there is an objectively real overlay of uncertainty veiling the phenomenon and makes the direct perception of its causality infeasible. It is this membrane which is a good image for the physical correlate of objective ignorance.

All this goes back to the Law of Continuity. The sum of a continuous function is a continuous function. In the same way, before we discover it, the proximate unknown is continuous with the ultimate unknown. Similar to Kant's

transcendental idealism in which the mind organizes objects into a spatiotemporal framework, like a grid knowledge superimposes onto ignorance and by that act gives the appearance that there is this or that unknown distinguished from all of the unknown.

The goal here is not to dig into the physics, but to use it to get across the point that without a means to discretely measure ignorance in terms of information, whether classical or quantum, there is no basis for concluding that it must be discrete. To measure is to find something out, to determine an information state. This leads us to revisit the indexing issue mentioned earlier "I didn't know *that*" does not index all that is unknown but only what appears to be a fragment. Thus its a case of identifying a known unknown, of drawing the grid. It picks out the one datum as a conscious object, an apparent single piece of ignorance. There are a few considerations to sort through here.

It's worth mentioning that all examples of negative ignorance also conform to Leibnitz's identity of indiscernibles principle—if there are no differences between a given pair of objects of comparison, they are the same thing, and self-identity is determined by an object tautologically sharing all of its own properties, including its spatiotemporal coordinates . No two instances of being (negatively) ignorant differ, furthermore no local instance of ignorance is discernible from global Ignorance in terms of its content. However, just as Tolstoy said that all happy families are happy the same way but no two unhappy families are alike, positive ignorance is highly variable. In the case of positive ignorance mental contents are recruited to sustain and extend it, and no two mental contents are exactly the same given that no two brain states, even in the same individual, are ever the same. Given that ignorance is nonphysical, though it supervenes on the physical, it is not subject to spatiotemporal self-identity, either.

A useful way of thinking does not have to be strictly true to be pragmatically valuable. The "that" which designates a seeming fragment of ignorance which has been recognized, flips it from a uniformly unknown unknown that is continuous with all that is unknown to a known unknown mixed state. knowledge, as awareness, indexes it, in a manner similar to quantum decoherence. This unknown has no definite referent until the utterer recognizes

what they didn't know and thus point to it. By doing so they entangle knowledge with ignorance. Ignorance is then caught in knowledge's pull and vice versa, and gradually, through processes of investigation, ignorance isdiminished in proportion to the growth of knowledge.

So to sum up, before knowledge steps in to identify a specific instance of proximate ignorance, that local unknown is undifferentiated from the totality of the unknown. So what is discrete here is the knowledge, not the ignorance. knowledge slices the unknown unknown up into manageable chunks, but the cuts it makes are projections of the mind, an imaginary dotted line we draw around the unknown unknown. This links what appears to be a fragment of ignorance to knowledge. Deep ignorance, unknown unknownness, has no discernible boundaries, and when we recognize a facet of it, it is no longer deep ignorance but has gotten tethered to knowledge.

Like a dark field at night, there are no edges to Ignorance. It is a seamless unity of nothing in particular and in general all at once, like space. Once we flash a light on a given spot of a dark field, we reveal what is there and thus immediately illuminate and isolate that formerly occluded fact from the surrounding darkness. That spot of light however is then no longer part of the darkness. It is useful to speak of proximate ignorances such as "he doesn't know how to cook or how to fix an engine" or "I have much more to learn" but these statements show knowledge. They parse known unknowns. The moment we determine knowledge of what is unknown, it is conditioned by the awareness we throw on it like the flashlight illuminating a spot on the field. It stops being pure ignorance. Intentionally objectifying a region of ignorance as a discrete fact is like measuring a qubit. At first it's all over the place until it is inspected, then it adopts a discrete state.

One major problem with this view of the continuity and homogeneity of ignorance is that there is relative ignorance. What is known and unknown is unevenly distributed between people. Not everybody knows the same things and not everybody is ignorant of the same things, although in some regions of Ignorance, most people don't know the same things, at least in a specific moment in time. (Do other planets contain life? Which ones?). If what is unknown varies

between individuals, then there can be no universal set of all Ignorance because for some agents a number of those elements will be known and to others unknown. Thus existentially quantifying ignorance breaks down. Speaking of a universal set of Ignorance is misleading because the assortment of elements included in this set must vary between individual to individual.

Such relativistic issues do not jeopardize the concept. Ignorance is analogous to the distribution of energy or dark matter. It does not have to be uniformly distributed, there can be more of it here and less of it there, but at the same time, all ignorance is self-similar all the way down. Energy, no matter the amount of it, is always energy all the way down. It has nothing else in it, and a large amount of energy is the same as a small amount or some fraction of its total with the exception of its magnitude.

Two people can know and not know the same fact respectively. However, knowledge is mind-dependent while ignorance is mind-independent, so the ignorance in question is in this confused in-between state where it is neither or both, divided between someone who knows the fact and someone who doesn't. The solution is that what is objective has to be evaluated objectively. What is unknown relatively then, is a kind of frequency distribution. That different people can have differing sets of knowns and unknowns, sometimes overlapping and sometimes incongruous, reflects how that ignorance, like heat, is distributed. One can picture it this way. The more people find out about a fact, the more the objective ignorance about it is exhausted, the more it shrinks toward a minimum value to the point it almost goes out of existence. If only one person knows about it, its decrease is marginal, there is a sharp deformation of the topology of ignorance collapsing on that person, but it is uniform elsewhere. Its contour changes, but it still retains a shape. If somehow everyone were to come to know the fact, it would approach zero. If the situation were to change, its value might fluctuate back up again, it would return to perfect homogeneity. It is dynamic and relativistic rather than absolute, like Einstein's relativity. So the fact that a region of ignorance is relatively known or unknown just means that its shape has been altered, collapsing on one person who knows here, opened up still to someone who doesn't there.

To clarify this last point with an analogy. Suppose several people for whatever reason have a large rubber band stretched tautly between them. Keeping each side of the rubber band taut is a pin that it has been stretched around. When someone decides to remove the pin, the elastic snaps into a different shape, but it still contains an area, and its shape is still maintained by the remaining pins. In the same way, suppose that the shape of the elastic band represents the topology of ignorance, as it is distributed among a group of people. Removing your pin does not completely break the tension of the band, it snaps it into a shape that fits around the remaining pins. To remove a pin is analogous to "coming to know", a moment of realization. It reshapes the contours of ignorance as it is distributed relatively among the group by freeing you from it, but its topological conformation is merely deformed, not destroyed. This analogy doesn't quite capture everything, but it gets the main point across. To know what another does not is to reshape the distribution of the ignorance that once applied to both.

Later, when we turn attention toward the theory of inquiry, we will indeed see that such apparent discretization of ignorance, even if it only gives just the appearance of fragmentation, is a key driver of learning, questioning and investigation. Appearances alone can be useful. To ask a question is a perfect example of how the mind can target a perceived localized region of Ignorance, fix it in place, and draw the appropriate information necessary to zero it out. It is impossible to throw ourselves at deep Ignorance all at once in its totality. We must, like a searchlight, scope out regions of its continuous surface and rope it off from the whole in order to fix attention and learn something.

> It is a common sentence that knowledge is power; but who hath duly considered or set forth the power of ignorance? Knowledge slowly builds up what ignorance in an hour pulls down. -George Eliot

Ignorance can be causal. It can make things happen. At first glance it's not entirely obvious how a negative concept can trigger an event. It's easy enough to understand how the prolonged lack of food is what causes hunger. On closer inspection it would be more accurate to say that the positive fact of the body's

unsatisfied requirements for sustenance results in the the activation of metabolic signals that cause hunger . Someone who is well fed can go on in absence of food until those requirements are once again unmet and the body signals hunger again. Similarly, there are certain drugs which suppress appetite, proving that it is not the absence of food which causes hunger but rather chemical messages that can be otherwise manipulated. By the opposite token, the consumption of food eliminates hunger when hungry: *if* food is eaten, *then* no more hunger. These types of conditional statements are the hallmark of causal reasoning.

Similarly, disease isn't caused by lack of good health, it is caused by the degradation of health. When negative concepts are causal, it usually means something is missing and this distorts how something behaves. Causal absence only makes sense with reference to a degraded positive fact or interrupted interaction between objects , a body whose normal function deteriorates into starvation or sickness. Nothing can come from nothing; negative causality is always about how a negative relates to a positive, either by preventing, skewing, distorting, blunting, diverting, diminishing, depriving, replacing, corrupting, its normative functioning and capabilities.

On balance the propositional form "X does not know Y. If X knew Y they wouldn't have done Z" and its logical inverse, "X knows Y. If they hadn't known Y they would have done Z" applies to many situations. This is generally the form of a mistake. If ignorance leads to a diminishment in cognitive possibilities, then ignorance circumscribes what options remain to be chosen, or it presents the wrong options.

One can imagine a series of branching rooms. Each room has a pair of doors on the other side, leading to separate rooms that in turn have two doors leading to further rooms. In each case, one of the doors is locked. In this analogy, the locked door represents an option about which one is ignorant. (By extension, one is therefore also ignorant of all the branches of rooms, both locked and unlocked, that stem from beyond that locked door.) To proceed, only those doors which are unlocked-which one knows- can be opened, and only the rooms they lead to can be accessed. In this example one can say that the agent goes with the door they know and can open, or they go with the door which is remaining after the one

they don't know and can't open is excluded. *These statements are logically equivalent.*

In the case of the rooms with the doors analogy, ignorance has a funneling or channelling effect. It is an overall structural determinant of the remaining options for knowledge. (One could extend the number of doors arbitrarily to elaborate on the point; only two doors where chosen in this example for simplicity.) Let's extend the analogy so that each door has a certain number printed on it. In each room there is a key with a number on it that can unlock the corresponding door somewhere in the tree. Finding such a key would be analogous to learning or rather *priming*, it is the ability to open up a further pathway into a different area of epistemic space because one has acquired an awareness that a certain action can be taken given a certain cue. Even in that newly opened up space unlocked by the key, however, one is still subject to conditions of knowledge and ignorance, the same sprawl of branching pathways. What one doesn't know has only merely been pushed back a level, and many doors are still locked. There is no master key that unlocks all the doors. Opening the door and making a choice means leaving behind all the other pathways, with different keys and different doors beyond it. Learning does not dispel the sum total of ignorance, it merely opens up new conditions where different cognitive possibilities are at play and more options are available, until they are not. Ignorance is also the result of an epistemic opportunity cost. We learn one thing at a given time at the expensive of another.

As a side note, one could also flip the door analogy around and imagine a situation in which rather than too few unlocked doors, there were *too many*. Ignorance would apply in this situation as well. Overabundant options leads to a selection problem where it becomes interminably undecided what the best option is to take. And since there are now exponentially more doors to keep track of, finding the right one becomes more difficult. Both too little and too much information leads to ignorance, the key is having the right information.

It's a universal truth that the sum of ignorance is always larger than the sum of knowledge, and that knowledge unfolds within ignorance. This is true invariably of who the knower is in question and is true over all knowers taken as

one. Here this obviously true statement will just be asserted, although later the claim will be more rigorously elaborated. In brief however, the mere fact of learning, the continually demonstrated capacity to find things out, furnishes ample evidence of the ultimacy of ignorance. There is always more to learn, and learning represents an iterative expansion of knowledge into ignorance. Every time we acknowledge we have learned something, which is nearly every waking moment of life, we reconfirm that there is more ignorance than was our previous knowledge. Moreover anyone who has stopped learning, refuses to learn, or claims omniscience, is ignorant by at least one more unit more than what they know. The inference that there will always be more to find out is the kind of inductive leap which is guaranteed to alway land.

In retrospect, noticing how truly we misjudged the adequacy of our past understanding when the future inevitably assails us with surprises is inescapable outside of the most egregious cases of self-delusion. Anyone who can sincerely remember what they expected of the world today ten years ago can measure the chasm between expectation and reality, and generally speaking, the longer the span of time, the greater this divide. One can conduct a simple experiment. Write down 3 major news headlines that will happen on this date of a decade from now . What are the odds of being right about a single one of them? History in hindsight often appears to be a comedy of errors, filled with misconceptions and drastically inaccurate prophecies that led to blunder after blunder. For every insight and intellectual breakthrough, there were medievals who thought the black plague was spread by evil winds, or Napoleons who chose to invade Russia in winter, or misguided physicians who mocked the idea that washing hands and sterilizing medical equipment could prevent infections. Or 1950s futurists who dreamed of flying cars and personal robots in the year 2000. It's quite typical to be wrong.

More will be said about fascinating properties of the temporality of ignorance. The point here however is that the vast negative space of ignorance enveloping knowledge, and often seeping into and distorting knowledge and belief, is no mere inert backdrop. It's a force of nature, a power unto itself. Put another way, it is much more probable to labor under ignorance than it is to labor under full

knowledge. knowledge is intrinsically a low probability state. To the extent that learning is in play ignorance is more like an organic *substrate*, defined as "the surface or material on or from which an organism lives, grows, or derives nourishment." Though learning knowledge eats ignorance, and as knowers we thrive the more we are primed toward capitalizing and "mining' the surrounding ignorance-as-potential-knowledge that our knowledge consumes. To the extent that ignorance is potential knowledge, to the degree in other words ignorance is negative, knowledge draws on it, is oriented and moves towards it, and gains order and complexity by reducing.

To tie this back in with the idea of causal ignorance. In the theory of general relativity within physics, space is conceived not as a passive vacuum or arena in which matter resides, but as a substance in its own right. Mass interacts with space, warping it. Sufficiently dense concentrated mass curves space, and it is this curvature which causes things to fall towards the local center of the densest nearby mass. This is gravity. There isn't in other words, a force of gravity in the sense of a positive influence tugging at masses through invisible lines of force and mysterious action at a distance. The involution of space, the substance in which mass embeds, channels masses into gravitational wells which they slide into. The effect of structural, causal ignorance on knowledge is analogous to this interaction between space and mass. Ignorance is what in mathematics is called an *attractor*, a set of values within a phase space (the set of possible values within a multidimensional system), towards which other values falling within its *basin of attraction* approach without never quite reaching (that is, asymptotically). Learning, in other words, is like falling into a local minimum of ignorance. It is that ignorance, the depression in the conformation of the system, which is the driver of epistemic action. So it's fair to conclude that negative ignorance *causes learning*.

A skeptic might complain that these are just metaphors, perhaps evocative ones, but still no more than clever images. Like the other physical analogies employed elsewhere, the curved spacetime analogy however does more work than that. It illustrates how what's normally conceived as a negative, recessive, empty element (space/ignorance) can interact meaningful with a positive,

assertive, contentful element (mass/knowledge). Is there a deeper connection between these analogs or are they merely illustrative? Perhaps not, but it goes to show that such causal relationships between positive and negative factors are not unheard of. Ignorance really does function as a kind of a recess that draws learners to it in their effort to maximize the mutual information between their internal states and external states until their learning goal is achieved. One does not experience curiosity in a setting where there is noting unusual, where there is no prospect of discovery. Bordom rather than curiosity would more commonly be the result otherwise. An interesting side-effect of this phenomenon is that boredom, like curiosity, is an epistemic or in this case anti-epistemic emotion. To become bored means to reach a state of inertia and insensitivity toward the surrounding distribution of ignorance. There are no "pockets" of the unknown worth exploring. One has become inured to the familiar and no longer perceives the novelty of ignorance inviting discovery at every turn.

Structural ignorance is different from the kind of ignorance with which we are most familiar in that it is more neutral. It is everything outside what is known, every alternative, counterfactual, or alternate universe from the one which we presently occupy, that we're not aware of. It is information entropy. It's not necessarily the kind of ignorance that is in principle bad or error inducing, though it can be. It is ignorance with a purely logical character, a manifold which is conceptually related to what David Lewis called possibilia or possible objects within possible worlds. If to know some things is to be aquatinted with it, as Russel observed, to know is to experience something. But every decision to experience is to take one route in an forking branching tree of possibilities. To make a decision is to select one path of experience over another, which is logically equivalent to declining to choose any other path. That is to say, affirming an option is logically equivalent to denying every other, it is disjunctive. To experience one thing is to not experience and thus to be ignorant of all others.

By excluding alternative courses of action ignorance can cause events. It structures what is knowable and conditions decision-making. It helps induce learning. It is the critical missing information which leads to a chain of events

unfolding differently than how it would otherwise. A lie that prompts someone to misrepresent the world and to behave with respect to that falsified version of events in mind. Mistakes are the consequence of a lack of knowledge and can be the cause of anything that human action can do. On a more upbeat note, ignorance also drives curiosity, for the simple fact that this experience would not occur in the absence of it.

> There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy (Hamlet, I.5)

A central concept to the vision of ignorance here described is *epistemological* or, somewhat interchangeably, *informational incompleteness*.[5] This concept speaks to the fact that all knowledge is necessarily partial, parenthetical, and subject to revision or updates, and that there is a lopsided imbalance between what one knows and doesn't know that typically favors the latter. The internal information state of an epistemic agent is eternally less than the information state of the world at large. This straightforward quantitative difference what that unifies knowledge, inquiry, and ignorance under a single universal principle.

All knowledge is necessarily unfinished. First of all, how can it be known that knowledge is necessarily imperfect, and one might say with a wag, how can that be known completely? It's perfectly plausible that a discrete set of facts as a *closed system with respect to itself* can be known to completion, but like beliefs, knowledge forms into systems and has a relational character to the world at large. Any set of facts, no matter how bounded with respect to what it references, is always relevant to facts that are beyond its scope. Both knowledge and belief are ultimately related to how the mind relates to the world, a world which is mostly unknown and about which ignorance prevails. As a system, knowledge is

[5] The difference between epistemic and information incompleteness is subtle. As will be elaborated elsewhere, information is the only thing that can be said to be "acquired" from the outside. knowledge acquisition its a misnomer, because knowledge is constructed from such external information signals, but can never be said to be signaled from the outside.

bound by its relations, and it more or less derives those relations from mapping the relations between facts in the world. But it is plainly obvious that any finite individual or collection of individual's experience of the world is subtends all possible experiences and facts afforded by the world. And in any case, every time we learn something new we achieve confirmation of this fact and proof of the incompleteness of our knowledge. In addition to these considerations there is William James' notion of an *unfinished universe*, a universe that is forever growing at its edges as it extends into the future. If the universe is incomplete because it expands into the future at its leading edge, it is hard to see how knowledge about could be anything but incomplete with respect to time and space.

Consider conceptualizing knowledge as a graph consisting of connections between nodes. Such a mental model of knowledge is an approach used in artificial intelligence. For the purposes of this merely illustrative example these nodes and their connections do not need precise definitions, but let us suppose that each node represents a factual proposition and each connection represents the material implication of one proposition by another. Suppose also that this graph exists within a manifold of such nodes, only some of which are connected to and associated by the graph. These "unassociated" nodes having no connections to any of those on the graph are analogous to ignorance as potential knowledge. Perhaps there is also a transformation rule by which the graph can "learn" by extending a connection to one of these disassociated nodes and incorporate it into the graph, but by the very performance of this act it confirms the fact of the graph's prior insufficiency. So long as the number of unconnected nodes exceeds those that are, the knowledge is incomplete. This little example helps show how knowledge is necessarily limited if the number of facts at large is greater than the number of facts it associates by the relations of knowledge. Assume also that as time passes, new unassimilated facts are added at the far end of the grid with periodic regularity as "current events" develop. So there is always a leading edge beyond the horizon of knowledge which implies the process must continue indefinitely without reaching a final state of maximum knowledge. New hosts of things to be ignorant of are added every day.

So incompleteness is reflected by limitations of experience but also the fact

that the universe itself is incomplete. Yet they are also so because of how the agent is situated in time and has a finite relation to their environment. Our observations of the world are contingent and partial because we are, to use Heidegger's term, *thrown* into the universe as it were *in medias res*. And the universe has existed for much longer than us and will continue to exist for far longer. It throws off so much information that we are always playing catch up. But we will never catch up. To top it off it exudes much more information each moment, which reflexively outdates previous findings and redefines them in a new light. Then there is knowledge of the future, which is drastically restricted because there is no information coming in from the future. The universe's ultimate beginnings and ends, its first and final causes, are broken off and fall away from us in time and we are confined to a small corner of the universe from which we are forced to extrapolate. We have to backwardly and forwardly reconstruct the fragments of information we collect from our finite standpoint relative to past and future. The unknown rushes at us and we must trace out a path of knowledge through it like a worm tunneling through the earth, destined to never know what is to our right and left and above or below but only what is right in front of us.[6] This incompleteness and the drive to resolve it is directly related to the problem of induction, a key concept to which deserves a separate treatment that reinterprets it in light of the theory of ignorance. Through induction smaller pieces of experiential knowledge are projected into larger maps that structure what it is expected from as of yet unexperienced, but still expected of objects similar to those which have been experienced. Ultimately there is a limit to how much we can leverage finite knowledge to encompass and anticipate the vast field of ignorance relative to which it unfolds.

A possible exception to the postulate of eternally incomplete knowledge is to be found in Hume's so called Principle of the Uniformity of Nature, which states that the laws of nature operate everywhere with undeviating regularity, and similar categorical generalizations. Such potent generalizations are indeed a

[6] In this case I refer more to experimental knowledge rather than conceptual knowledge . i.e, knowledge derived from the senses rather than by reason. Generalizations which transcend experience tend to go beyond the reach of it.

special class of knowledge. However, there is a distinction between the knowledge of a rule and the application of a rule. I, or a specialist physicist, may understand gravity to an adequate degree, but this does not tell me about every situation to which gravity applies. Such generalizations are still subject to the relational nature of knowledge. knowledge of a law of nature does not translate to knowledge of how such a law interacts with the entirety of the facts it subsumes under its principle.

The generalization of inferences from an incomplete sample of instances to laws of nature, to explain the type in terms of the tokens, is an inescapable byproduct of incompleteness. Inductive generalization are therefore like a lasso or abstract "dotted line" drawn around a class of both experienced and unexperienced objects projected to share the same defining class properties. Generalized knowledge is still porous and full of holes, it merely demarcates a zone of phenomena about which we are ignorant in the specific but which we we can assert that if observed and discovered will conform to the general rule. Although the most useful forms of knowledge typically take on the compact economic form of valid generalizations, it is precisely the generality of this form which obscures how much we don't know about the application of the rule to each instance. On the flip side, the most deadly forms of ignorance are bad, hasty generalizations, positive ignorance that impostures as systematic knowledge. As such they misrepresent greater sections of reality as they gain in generality like an irrational witch hunt that adds more and more to the list of the falsely accused because of faulty and loosely defined criteria for guilt.

Incompleteness is also the intrinsic rationale behind learning. The capacity to learn is an inbuilt design feature for accommodating information incompleteness. All learning is proof of the fact. As previously stated learners are driven towards what they do not know, and this orients and powers their drive to understand. As finite systems, minds are positioned before a vast field where more information resides outside the boundary than within it. From this elementary, brute fact much can be inferred. Information incompleteness and the learning it drives is an attempt to rectify an inequality between outside and inside information so that the agent is brought into a more optimal alignment

with its environment. When someone moves to a new city, they walk around to get the lay of the land. Typically they do this until some subjective satisfaction criteria are met rather than until they absolutely memorize every last detail of the city plan. To read a book is to reach an level of parity between its contents and one's understanding of it. In any case, all learning moves towards a greater degree of equilibrium with external information. Through extensive study we may exhaust all the unknowns of a local region of ignorance but in absolute terms we have only drunk a glass of water but not the ocean.

In what follows we will examine the many ways that epistemic incompleteness defines a fundamental relationship between the (un)knower and the world.

The relationship between knowledge and ignorance can be loosely likened to the relationship between matter and dark matter in physics. Weirdly, in a way most of what is isn't. The universe mostly consists of an unidentified form of matter unlike the typical baryonic matter which makes up familiar, perceivable stuff. The universe is thought to be made up of 85% of a material that is unobservable using conventional instruments—and antithetical to the principles governing observable matter. This strange stuff must exist to balance out the cosmological equations and to explain phenomena such as cosmic expansion. Like structural ignorance, dark matter must play a highly relevant role in conditioning baryonic matter just as ignorance conditions knowledge.

Dark matter does not absorb, emit or reflect particles of light. It cannot therefore be detected by conventional means. knowledge is often metaphorically conceptualized as light. It is enlightening, illuminating, it casts over the darkness, which oppositely is associated with ignorance. We speak of people who are intelligent as bright, and exceptionally informed people as luminaries- literally, light-emitting bodies. Meanwhile people who are ignorant are benighted, dim, in the dark. How then to make sense of dark matter, this mysterious substance which is but isn't, which cannot be detected by normal means but still has a role to play in the cosmos and is utterly unlike the matter as is normally understood? Of course, this metaphor about dark matter breaks down when we consider that

what much that is shrouded in ignorance can be known. This is not meant to be a perfect analogy but to underscore the objectivity and substantiality of ignorance.

The origins of the light/dark knowledge/ignorance analogy are quite plain. Human beings are visual creatures, and so light is one of the primary means by which they gain information about the environment. Light quite literally is a source of knowledge. Equally, for the opposite reasons darkness is a source of ignorance. This eternal Manichaean distinction between light and dark, knowledge and ignorance, is cliche precisely because it stands for an actual universal experience.

The dead space of the unknown shadows the material that is present to knowledge. In a related way, what is unknown exerts a kind of gravitational pull on what is known. To know is always to be drawn towards and to seek out where that knowledge leads and ends. To notice that some form of ignorance is impeding you from achieving a goal is to drive headlong into it, to negate a negation or subtract a deficit until the amount of ignorance approaches zero or the goal is achieved. (At least when we behave in a cognitively productive and rational way.) Ignorance pulls us forward even though we can't help but walk into it, constantly bumping into things we didn't know and being caught by surprise. The frontiers of life are marked out by such surprises, by the fringe outer perimeter of what we've discovered through a partially planned, partially random walk through the all-encompassing space of ignorance. By contrast with it, we make sense of everything that falls within these outer boundaries and call that the known universe. There are seemingly infinite things we could know. Like dark matter, ignorance is native, while normal matter and knowledge is just a tourist. Again, knowledge is the exception and ignorance is the rule.

Given the supremacy of ignorance and randomized chaos and the observation that we are guests of its cosmic casino playing by its house rules, and given the comparative smallness of ordered knowledge compared to it, it is strange that ignorance has not gotten its due in the philosophical literature until fairly recently. knowledge unfolds within the ignorance, and just like how the dimensions of an object are defined by the edge at which it stops and the space that continues beyond it, or a beam of light can be made out clearly in

surrounding darkness, so too is knowledge shaped by it.

> When we measure something we are forcing an undetermined, undefined world to assume an experimental value. We are not measuring the world, we are creating it. -Niels Bohr

Kant introduced a distinction between what he called the phenomenon and the noumenon which has had an enduring influence on epistemology and ontology. What we are able to directly perceive and experience, is phenomenon. The phenomenal content of perception is structured by our sensory and cognitive schemas that are trans-empirical: the mind formats experience with its preconfigured representations. Consequently there is no way to be sure that our perception of the world is how the world is in itself. We experience what is not a pure reproduction of things as they are in themselves but instead how the mind constructs it. Thus there is also noumenon, the "thing in itself", which is the underlying substrate of objective reality that exists independently of perception. In other words, the object of the senses which is the source of knowledge, the phenomenon, is in a way adulterated. It has the categories of the understanding imposed on it, becoming something other than what it corresponds to in the external world. Just as how color is a property of the eye that sees and not the object seen, as witnessing the object under various lighting, or afflicted by color blindness, shows, so too is the inner nature of an observable other that our perception of it. The object as it exists in itself unperceived has a nature that is unknowable. Thus, Kant concludes, while we can know the noumenon exists, because sensation is purely receptive and therefore must be caused by something objective that is outside perception, we cannot know its nature directly as it exists within itself. The noumenon is unknowable.

I wish to suggest a subtle variation on Kant's conclusion. Given the substantiality of ignorance, the unknowable *in itself* shares a certain resemblance to Kant's noumenon. It is not that the thing in itself is unknowable. It is *the* unknowable. That is to say, the unknown is the primary, default fact of reality, it

is the "something" that precedes perception. This slight grammatical emendation of adding the definite article "the" to Kant's conclusion transforms the unknown from an adjectival description, "is unknowable" to a noun. It turns the copula from an ascription of a property to the assertion of identity between two things. The unknown, and the subset of the unknown that is unknowable, is a reality, a self-standing fact, although it is not a thing in the conventional sense as having spatiotemporal extension. It is not a thing among things, but it has a categorical reality. It is analogous to a fundamental field in physics that pervades the universe and of which observable physical particles (knowns) are "higher energy" excitations of it.

The immediate rebuke to this identification of the thing in itself with the unknown is that the unknown cannot be the thing in itself because we can come to know the unknown. However, as soon as the unknown transitions to known, it does not somehow retain this identity in the same way that the object that is perceived is not the same as its noumenal substrate . The unknown merely recedes back, it is as it were never "touched" by knowledge but recoils as soon as it is replaced by it. knowledge displaces it and like a weight placed in a bucket of water, moving the surrounding fluid around it without actually eliminating it.

The more familiar internal psychological ignorance is counterposed to an external, substantial and physical ignorance which is embodied by entropy and the raw, unprocessed, informational noise of physical reality. Though the physical processes that ultimately generate ignorance can be understood, the ignorance they "cast off" is ontologically prior to them. They are surface expression of an essentially epistemic (informational) noumenon.

Granted, a skeptic might reasonably feel this assertion is farfetched . But this reluctance to accept the deep objectivity of ignorance stems from engrained but faulty intuitions that ignorance always exists relative to a subject's mental states. It recalls physicist John Wheeler's formulation It from Bit. The "It" of phenomenal representation (knowledge) is the "Bit" of the underlying ontological substrate (ignorance-as-uncertainty). Given that bits are measures of uncertainty, and uncertainty is a measure of ignorance, the logical consequence of this line of thinking is that the ultimate ground of being is ignorance

appropriately defined. As Wheeler writes (pg 312),

> It from Bit symbolizes the idea that every item of the physical world has at bottom — at a very deep bottom, in most instances — an immaterial source and explanation; that what we call reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin and this is a participatory universe.

Information, or rather the uncertainty which it measures, on balance, is a good candidate for the noumenon as the transcendental object. The way Wheeler describes it, information is the hidden, fundamental layer of reality which surfaces as the probabilistic wave function and other metrics of quantum mechanics. You cannot sense, see, hear, touch, taste, information "in itself" in a pure form that is detached from a sensory correlate. To sense is still to register information, to close a loop. Speaking of photodetector experimental equipment which increments a counter when a photon hits it, Wheeler points out that the counter clicking is the only proof of the photon's existence. Specifically with regards to quantum physics in this case, the photon prior to measurement does not exist in a determinate sense but rather in a superposition of all its states, or equivalently, a spread out probablistic superstate. Only once it hits the detector can a definite causal history be traced from it and we observe that the photon took a given path. Analogously, knowledge is derived from ignorance. In his paper Wheeler counters the claim that existence thus built on "insubstantial nothingness" is nonsensical. This informational nothingness is meaningful. After all, he points out, Rutherford and Bohr made it clear that a table is 99.99% emptiness, with the distance between the nucleus and its orbiting electrons proportional to that of a golfball and a speck of dust a kilometer away. Yet the table is there, informed by an underlying "nothingness" that is still a perceivable something.

Information does not exist on its own in a pure state. How then to countenance the seeming contradiction that information is fundamental but cannot exist without something that it references? This would seem to reverse the script, suggesting the more common sense view that information is indeed

secondary, “Bit from It.” To escape the infinite regress of having the universe rest on some underlying foundation which is not self-standing (no turtles all the way down), Wheeler posits a loop, what he causes the participatory universe. According to this dual aspected theory, information is finalized when subject and object connect through observation, and the insubstantial unknown becomes the substantial known. As potential, the unknown does not have meaning separately from the knower even though takes ontological and epistemological precedence. It is one side of the same coin, and you can’t have one without the other. By this same token, epistemology *is* ontology because information *is* physical. In a participatory universe you cannot have one without the other. Or as Wheeler put it, “physics gives rise to observer-participancy; observer-participancy gives rise to information; and information gives rise to physics.” (p. 314)

This notion of a participatory universe is also connected to the theory of inquiry. Regarding the issue of quantum measurement, the type of device used determines the outcome: “the choice of question asked, and choice of when it's asked, play a part — not the whole part, but a part — in deciding what we have the right to say,” (p. 316) says Wheeler. On the quantum level, information is not “out there”, it has to be retroactively identified by the process of measurement. Thus “no question, no answer." Which is the same as saying, “no ignorance, no knowledge.” Analogously, the process of investigation in general contains the seed of its object. Inquiry points the way towards its selected unknown, targeting and detaching it from an amorphous background of pervasive undifferentiated ignorance and preparing it for discovery as a known quantity. The unknown unknown is bracketed and compartmentalized by inquiry, bit-flipped from unknown unknown to the known unknown, and drawn forward into clearer awareness as inquiry progresses. The intentional object of one’s own ignorance recedes in proportion to which knowledge is obtained. Once the former is minimized in relation to the latter’s approach to its maximum, the loop is closed. Investigation ceases in other words at the point at which one believes one’s ignorance of its object of inquiry has been exhausted or the object is in fact fully neutralized.

Through inquiry we are active participants in an epistemic universe the

ground condition of which is ignorance. Inquiry, and by extension perception, gains purchase into a pervasive, generic, undifferentiated, unbounded unknown and through its questions sharpens aspects of it into definite, answerable, bounded forms: known unknowns. On this point, perception, measurement, and inquiry each share an identical relationship to the thing-in-itself to which they interface. They are analogous to the photon counter clicking when the light quantum hits it; processes by which the indefinite is made definite and from which further perception, measurement, or inquiry proceeds from that information which has since been "anchored" by discovery. At which point, to cite Bohr, the indefinite state is "brought to a close" through an "irreversible act of amplification." (Bohr, 1935) Indeed, observation and measurement at the most basic level are acts of inquiry, probing behaviors that fall under its rules and logic. Any organism scanning its environment, or any act of scientific measurement, are ultimately bound by the same principle. When a physicist measures the bit registration of a photon, they are implementing the same logic of ignorance resolution in miniature which I argue is ubiquitous to all thinking processes.

The claim that uncertainty/ignorance is the thing-in-itself might come off as dissatisfying, paradoxical, nonsensical, or perhaps even nihilistic. This is the "insubstantial nothingness" which yawns before us, and it deflates any hopes for absolute knowledge that could banish all mystery . It seems to suggest an impossible vagueness is at the bottom of it all, in stark contrast to various religious or philosophical formulations that say in the beginning was the Word, the Logos of rational order and intelligibility that glues the universe together. It violates Leibnitz's principle of sufficient reason, the statement that everything must be rationally explicable. The universe then , Leibnitz suggested, must ultimately be perfectly knowable because it is informed by the mind of God, who is perfectly rational and all that is rational is knowable. These views are not incompatible however. Uncertainty and ignorance, properly understood, are not anarchic or senseless by any stretch. In their madness is a method. In the last, what matters is not ignorance, any more than what matters about a Michelangelo sculpture is the raw marble in the mountain it was taken from and carved out of.

Like the unobservable essence Kant supposed hid behind sensory data, ignorance stands behind the known world. It's what we make of this raw substance and how we orient ourselves toward it that matters.

Ignorance and Information Processing

Ignorance can be interpreted in many ways. Given that ignorance has an essential cognitive aspect , and given that information processing accounts is in vogue in cognitive psychology, it can be fruitfully analyzed in terms of information processing. When thinking one does not mull over knowledge, after all, if we have knowledge there is no reason to think: we just know and act off the basis of that knowledge. knowledge has an immediacy: it may take us time to recall knowledge we already have, but this memory retrieval is not the same as thinking, and in any case once retrieved the knowledge is immediately invoked. Thinking leads to knowledge. The goal of thinking is to arrive at knowledge, or at least the resolution of doubt, which is halfway towards this ultimate aim. Thinking is *knowledge formation*. Uncertainty reduction, and thus the pursuit of resolving our own ignorance, is the motive for cognition and the prerequisite for knowledge. In that respect thought is intrinsically negentropic and wound up with information entropy.

Thinking— rationally effective thought—can therefore be conceived as a process of disambiguation. It separates the wheat of truth from the chaff of falsehood. It is intrinsically a negentropic operation in that it sifts a disordered entropic mixed state in which the truth and falsehood are disarranged and interspersed and settles them into a defined gradient that clearly tells them apart. This association between entropy and cognition is a growing area of research, and much more will be said about it. For now a general overview of the idea is in order, because this model of thinking in relation to ignorance is central to everything else.

Of particular interest are the rigorous conceptual toolkit of information theory and the concrete models of artificial intelligence. These systems bring ideas to life, concretely implementing concepts as material demonstrations which would

otherwise be merely theoretical curios. They are actualized instances of ideas that actuate in the world of physical hardware, and because they work, there must be some truth to them.

In a similar vein, mathematician David Wolpert's work developing an impossibility proof of a universal inference engine, particularly the inference device on the scale of Pierre Simon de Laplace's famous Demon thought experiment, tells us that no physical machine that is part of the physical world it observes could ever achieve omniscience or perfect inferential salience Always there is more "in heaven and earth" for such an inference engine to deal with and so their inferences are capped at an upper limit. This is an important finding for arguments about unknowability claims. While AI programs are not perfectly inferentially valid models of the human mind, they can nevertheless be thought of as isolated pockets of structured information processing that implement problems that it was once thought only humans could solve. Even for artifactual programs such a AI it is realistic to speak of ignorance-as-uncertainty constraining and conditioning their computations just as it does our thoughts. This establishes more evidence that information incompleteness is a universal phenomenon affecting anything that must process information, which in later chapters we will even extend to systems which are not even conventionally thought of as cognitive in disposition, such as evolution and the universe itself.

The perspective on ignorance put forward here sees uncertainty as centrally important. To whatever extent that ignorance can be quantified, which again is only in relation to how much awareness there is of it, uncertainty is a factor. Uncertainty is a primitive relation of ignorance but is more directly tied to information processing. It can be more precisely tied to information than knowledge. It is quantifiable: in theory, the measure of uncertainty should be the sum of how many indeterminate random variables, or unanswered yes or no questions, considered as binary bits are required to fully characterize the probability space. To take the simplest, oft used example, the uncertainty of a coin toss is .5, so the information content of a coin is 1 bit given that there are just two alternatives.

This technical, quantitative approach to uncertainty is difficult to graft onto ignorance and knowledge as higher level concepts. What is the measure of my ignorance of who will be the first person who will step foot on Mars? It would seem that my ignorance of this question goes beyond measurement, because I am unable to even list a set of candidates more precise than anyone There is no definable "probability space" encapsulating it. Or I can say my uncertainty is absolute. The obvious response to this is that to measure something you need something to measure, there needs to be a signal, a spike of difference, that sets some candidates apart from the backdrop. Suppose that we were on the eve of the Mars landing, and there were 6 astronauts aboard the ship. The information could then be measured at a little more than 2.5 bits. Suppose the names of the astronauts were Ned, Ted, Fred, Jill, Bill and Phill. First you would ask, is it Ned, Ted or Fred? (Y/N), if yes, you would then ask, is it Ned (Y/N) or is it Ted (Y/N). Or if the answer to the first question was no, you'd do the same for the remaining 3 astronauts. This would mean that there is an average of about 3 questions in 4/6 cases, 2 questions and 2/6 cases, which averages to to the need to ask 2.5 yes or no questions. Perhaps in some weird way that goes beyond human understanding I could calculate a similar decision tree for all people on earth and ultimately figure it out. As it stands in the absence of all evidence the probability of whoever will be the first to walk on Mars is the same for everyone, so the uncertainty is the maximum. You might get pedantic and rule out everyone who would be too young or too old, or who would not pass the physicals or who was not alive at the time the Mars landing happens. That however goes far beyond the point behind made.

Uncertainty is not present for all degrees of ignorance. One is not uncertain about what one doesn't know one doesn't know. Neither is one certain in this case— there is only obliviousness. There can be no uncertainty where there is no awareness. And one who is positively ignorant is sure of their false convictions. (In this situation uncertainty could be a path to knowledge, a moment of clarity that relaxes stubborn commitments to falsehoods.) The relationship between uncertainty and ignorance is therefore slippery. While you cannot have uncertainty without ignorance, you can have ignorance without uncertainty.

Though the experience of uncertainty is a knowledgeable recognition of ignorance. By the same token, the role of certainty in modulating epistemic commitments to beliefs is also thorny. If one is certain in a false belief, one is under the spell of positive ignorance. In this case, negative ignorance (the lack of knowledge that one is mistaken to believe the false proposition) misestimates the confidence-evaluation of what's true. Conversely, the certainty that one doesn't know something is an inversion of this relationship, the truthful realization of one's ignorance as to a fact. All the same one must be certain one knows something to truly know it under positive introspection, i.e to know one knows. Else this incidentally factual conjecture is more like a lucky guess than a valid inference generated from a model of the world that is well adjusted to reality. Ultimately what is of greater importance is not the assertion of what is true itself but the validity of the model or theory that generated it, as this does a much better job of explaining the how, why and what of a fact rather than that it just is. If one doubts a model or theory, from a pragmatist standpoint that means one will be hesitant to use it, and so knowing that one knows always presupposes certainty.

It can most accurately be said of both certainty and uncertainty that it is expressed as an metacognitive evaluation of adequacy of one's state of knowledge. It is an expression of the degree of epistemic confidence. What is uncertain is not what is unknown in itself, the fact itself that is out of reach. Uncertainty arises when the subject is undecided as to how much what they know is in comparison to how much they don't. This includes how much they know they know. Uncertainty is the *self-recognition of one's epistemic incompleteness.* Uncertainty is therefore tied to known-unknown interactions. For an agent to experience uncertainty is to be aware there are several possible outcomes rather than just a single necessary one that is fatalistically impressed upon them. Or, it is to recognize that, whatever the outcome, one is not confident in what present available information and understanding foretells . Thus uncertainty has both an objective and a subjective component, which is to say, one perceives it in terms of how many alternatives an event affords by its intrinsic nature plus how much one is confident one of these alternatives will

play out or whether there are others that are unidentified. It's associated with arbitrariness: if the overwhelming evidence suggests one outcome, then other nominal possibilities are discarded. Uncertainty is never without the possibility of an outcome about which one is aware. One always knows when one is uncertain, and usually one knows what one is uncertain about. Uncertainty therefore impresses itself on us as a force, whenever what one knows comes into friction with what one doesn't uncertainty sparks up.

Critically, according to Shannon's formulation uncertainty it taken as a prior , baseline background condition which prevails prior to the determination of the information value of a signal. It is what needs to bee sorted out to sift the signal from the noise. So the more uncertain a phenomenon the more information needed to determine it; the information is a *measure of how much there is to find out*. Uncertainty has a definite value because the amount of information needed to reduce it to reduce it to infer meaning is proportional to how much information is needed to distinguish the message from the mess. Because possibility is necessary for there to be probability, another way of saying this is that a maximally uncertain event is one where the maximum of possibilities is equally probable. This formulation is referred to as the principle of maximum entropy and it will play a signifiant explanatory role later. This notion of "narrowing down the possibilities" in the course of acquiring information is fundamental to the rendition of ignorance-based epistemology. It brings together possibility, probability, entropy, information, ignorance, and knowledge in the same conceptual framework.

Subjectively, the key feature of the logic of uncertainty is how it relates known and unknown by degrees, that is, through probability. One may know all possible outcomes but be uncertain which one will be realized. One can also know that just *something* will happen, without knowing at all what the options are. There can be a mix of known knowns and known unknowns to any evaluation of uncertainty. However there cannot be unknown unknowns. Instead, the very cause of uncertainty is the concern that there may be unknown unknowns about which one is unaware relevant to the situation. Such complete ignorances are blindsiding, coming from left field. Probability is the quantitative

method for parsing these relationships.

Though it plays an important role in information theory, surprise and uncertainty are not the same. It may be said that uncertainty is the expectation of surprise. Seeing that the information value of a datum is a function of how surprising it is—meaning , how improbable—a state of uncertainty is one in which one is maximumly exposed to surprise. To put it in a somewhat aggravatingly roundabout way, to be uncertain is to be certain that one will be surprised. Doubt, therefore, is a form of uncertainty that exposes oneself to the possibility of new information. Accurate certainty therefore is a state in which all information has been exhausted, a state of equilibrium between the agent and their information gradient.[7]

Arguably there is a deeper, more mysterious side of uncertainty that is embedded into the fabric of nature itself. This is tied to the uncertainty principle in physics, and to the question of the ontological status of pure randomness and strict indeterminism. It's difficult to prove the existence of a purely random phenomenon, such that no amount of information could ever rule out the possibility of being surprised. That's because it is not known whether it is in principle impossible to collect enough information to see that what we thought was random was really just the result of a broader deterministic non-random process the causes of which are unknown.

It is no contradiction for uncertainty to be dual aspected. Both cognitive evaluation and physical fuzziness contribute to the same conditions of uncertainty. In a physical universe lacking in any sort of "noise injection" shaking things up and keeping them from strictly repeating the same behavior,

---

[7] The topic of senescence brings up another intriguing scenario in which certainty begets ignorance. Incorrigible, undoubting senescence is tied to an intolerance for questioning one's premises. They have become so intimately woven into one's core assumptions and perceptions of the world that to revise them would prove too destabilizing. The psychophysical parallel to this is the physically correlated slow down of the formation of new neural connections which necessarily retards further learning. Certainty can be a sign of cognitive rigidity which blocks learning just a much as a source of confidence that improves a knowledge model through its application. The false certainty of the senescent mind therefore lies in thinking that because one is no longer responsive to new information there is nothing left to learn. Willful ignorance is a related quality though it must have a different etiology.

there would be minimal uncertainty to cognitively evaluate.

It fits well into the view of world presented to think of some events as being intrinsically uncertain and volatile. Their causality is disordered. This fuzzy causality impresses itself on a subject and it is this which they directly experience as their own ambivalence. Looking at it this way, some events *put uncertainty into* a subject, it is not that the subject approaches the definite causally ordered event with their own uncertainty about it. This is a twist on the familiar way of thinking about the uncertain. It's typically thought that we just need to clear our head of its fog and that the answer would then be definite and obvious, that in other words our uncertainty necessarily speaks to our own cognitive limitation. Once we stop getting in our own way by acquiring more information or by discounting false hypotheses or thinking through the problem correctly, then we see the truth on the truth that misunderstanding obscured. It is instead argued that sometimes what are considered to be misconceptions are actually the byproducts of an inherently uncertain situation. It may be the case that there are times when the world itself has not made its mind up, so it is impossible for us to do so. Truth may be pending, rather than determined all at once. The uncertain cognitive state is then a *perception* of what *is* rather than a *confused estimation* about discerning between what is and what isn't. We thus *sense* the uncertainty; it is not identified with the associated mental confusion that indicates the uncertain but the fact of uncertainty as it weaves into the causal processes of the world including our minds.

Such a claim about the nature of uncertainty obviously creeps into some very thorny metaphysical thickets. The best way to approach it is by digging deeper into what information theory and related topics can teach us about uncertainty and ignorance in general.

The thesis that ignorance is substantial leads to a provocative insight. The uncertainty that the doubting mind experiences is not only an attribute of their own undecided cognitive state but a extrinsic property of the world at large. Oddly, what they perceive at first is the real truth of the world as it is in itself and

the regularities they eventually see in it a kind of distortion: the unobserved exists in diffuse suspension, while the observed is a subjective order imposed on objective uncertainty. It is not the mind that is clouded, the world is clouded, and it is simply seen as is. Ignorance is after all normal, knowledge is exceptional. Being ignorant therefore is strangely being in tune with the ultimate truth of the epistemic universe, and possessing knowledge is a kind of defiance of it. Ignorance, not knowledge, mirrors the true state of nature. In its absence of representation it is representative. Put differently, the world is "blurry" in its fundamental nature.

From one angle the claim that someone who doesn't know is "seeing" the ignorance outside of them is indeed odd. We tend to think of knowledge and accurate perception as a kind of correspondence with the truth, a banishment of illusion. In knowledge the mind and world align. The mind finds the patterns that are actually there after sifting out what isn't. On the contrary. Whatever patterns the mind perceives it puts there or rather arrives at. In other words what the sense organs perceive is what they are structured to detect, not what is really there.

Perception is often taken as a kind of tacit knowledge. Although hallucinations and misperceptions do occur, these are seen as deviating from true perception, which conveys accurate information about the external world. This recalls also the argument that if perception lies, then it would not lead to a survivable organism because an organism whose senses are geared toward showing it imaginary predators will not aid its fitness. However these kinds of misperception are not what's at issue. Such a position fails to appreciate the utility of certain self-interested reality distortions. Biological perception itself is sculpted around the particular fitness requirements of a given species. In sensation, reality is put into the service of perception rather than the other way around. As writer Robert Anton Wilson once observed:

> Every kind of ignorance in the world all results from not realizing that our perceptions are gambles. We believe what we see and then we believe our interpretation of it, we don't even know we are making an interpretation most of

the time. We think this is reality.

While in actuality not all ignorance arises from such perceptual disruptions, subjective reality distortion stems from an inherent bias the tendency to loop everything back to the organism and its needs. What if, however, "the gambles" of perception aren't something we add to the world, but something innate to what is perceived? What if the world is intrinsically physically "blurry"?

Who has a more accurate perception of reality in the following scenario? Two people become aware about the same subject about which neither of them has experienced. One is certain that such is the case about it, the other is unsure what to believe. Nothing about the object fixes their perception of degrees of certainty because neither of them have seen it. Supposing that neither has any more information about the object than the other, the certainty of the convicted one reflects merely an arbitrary will to believe rather than a rationally calibrated expectation value. They are wrong to be certain. It is certainty which is a product of the mind. What is not observed and known exists in a superposition of possibilities. Like the proverbial Schrödinger's cat, the unobserved has an ambivalent unassigned truth value, its informational degrees of freedom are at maximum until observed. Indeed prior to observation, the certain person is assuredly less correct than the uncertain unless they happen to make a lucky guess. The uncertain person is just telling it like it is. Any regularities the object might display has yet to be impressed upon them, an in the absence of any evidence, the maximum of its undetermined possibilities is the true state of the object.

Suppose now however that one of these observers investigates the subject in question and finds out definite facts about it and the other still doesn't. Should not their perception of it reduce the uncertainty as they see possibilities become actualities and over time get a feel for what the object is? Which of them is correct in their perception? One knows, the other is oblivious, and the fact of the object cannot be both definite and indefinite simultaneously. Normally it would be seen that the one who has investigated the object has grown closer to reality by the fact. They have lifted the veil and found what lies behind it.

According to the problem of induction, to assume the object undergoes some form of change—and everything does eventually—means it is impossible to be sure of it what you know of it is fully true of it. The observer of the object may observe the object behave one way in each of 99 cases only for it behave differently in the 100th, and they cannot know it will repeat this for the 199th and 200th intervals without continuing to observe it in that many cases. Thus the one who hasn't seen the object is still right to be uncertain of its possible states, even if they lack more refined statistical priors about its actual manifest behaviors. Even, in fact, if they have less knowledge of the object.

Should the observer's perceiving of the object not affect the uncertainty of the other person, deciding the undecided once and for all? After all, they can tell the other person what they saw, and then the other person can confirm it for themself. This is not the case because, once they have found out it does not collapse all sides of the object's resting state of innate ambiguity. It rather collapses one side relative to them. If that one person is the only person in the world who is certain of the fact, they exist in profound asymmetry to all other observers. To all else the uncertainty remains between them and the object. News of the knower's discovery does not propagate to the rest of the world through the act alone.

Let's suppose that the object of uncertainty is a piece of paper with a letter written on one side placed face down. Someone must guess the letter on the reverse side. They guess and then turn the card over and see that it is the letter G. Now, the letter G was always on the side of the paper. The flip side of the card did not not cycle through all possible letters in a state of indeterminacy prior to being flipped. (Though something like the Schrödinger's cat thought experiment suggests that that.) It was a persistent if unknown fact that would repeat as many times as one flipped the card or as many times as new observers came to flip it. The essential factor in play here is not the letter on the face down side of the card, however, but the *binary information state* of the card itself as either flipped or unflipped. The state of this card is itself one bit of information, and what is objectively uncertain is not the letter on the card but the one bit of information which obscures what that letter is. The uncertainty about the letter is a

consequence of the uncertainty of the states of the card flip. To the observer who has yet to flip the card the letter could be anything, because *is* anything. In the face of the unknown all possibilities are equally true.

Once a fact about the world is known, that fact is registered in place and persists as encoded into the state of the world as much as memory. In the same way that the open-ended future, through collapsing into and funneled through the pinhole of the present, becomes at last the stable and cemented past which cannot change and cannot be causally influenced by future events. Consider Schrödinger's famous cat.

Briefly, Schrödinger's cat is an attempt to demonstrate the paradoxical nature of quantum superposition which apparently violate the law of the excluded middle. How can two mutually exclusive states exist simultaneously? A cat is placed in a box with a vial of poison gas which is connected to an highly sensitive particle detector that is aimed at a bit of radioactive material. Once one of the radioactive atoms decays, the detector picks it up and this causes it to release a hammer that smashes the poison vial, killing the cat. According to quantum physics it is impossible to determine prior to observing the system which of the states, dead or alive, the universe will "decide on." Though originally conceived as a purely analogical illustration of quantum behaviors, later work by Huge Everett III introduced the notion of a universal wave function, which essentially generalizes the principle of the microscopic state to the macroscopic state of the entire universe. Effectively what is true of the quantum is also true of classical objects at humanly perceivable scales. In which case the cat's condition would quite literally exist in all states simultaneously[8].

---

[8] More precisely, Everett, had a different interception and claimed that the wave function took all paths simultaneously, but that they branched off into separate universes, hence his notion of the multiverse. Thus the wave function conserves pure determinacy and under this interpretation there is no objective ambiguity. There is no mysterious superposition in which the ultimate state of the system is unknowable, there is a universe for each state. On more careful inspection this does not refute the argument presented, because from the

Following his original description of the cat experiment,[9] Schrödinger makes an intriguing comment:

> It is typical of these cases that an indeterminacy originally restricted to the atomic domain becomes transformed into macroscopic indeterminacy, which can then be resolved by direct observation. That prevents us from so naïvely accepting as valid a "blurred model" for representing reality. In itself, it would not embody anything unclear or contradictory. *There is a difference between a shaky or out-of-focus photograph and a snapshot of clouds and fog banks. (*Section 5*)* [emphasis added]

This last sentence is particularly evocative it its suggestion that uncertainty is both in the perceiver and the perceived and its clear distinction between subjective and objective uncertainty. Uncertainty wholly within the mind is the out-of-focus photograph, the noise blurring the representation of the distinct object is introduced by defects in the mechanism that captures it. A high fidelity snapshot of something inherently nebulous, however, will share the same semblance of blurriness though the source of it is the indistinct *object represented clearly*.

Schrödinger then proceeds to explicitly mention the epistemological implications of this line of thought, which he describes as an "about face." He wrestles what he calls the classical model versus the blurry model of physics. In the classical model, ultimately what the physicist comes to learn about the system, to the extent that its behavior can be predicted, is a complete description of the system. Prediction is proof of complete knowledge. This is an epistemology in which complete knowledge exhausts the system observed, and all parameters are determined by initial conditions. In a passage that echoes the notion that thought is a narrowing process, Schrödinger (section 1) states,

> [By refining hypotheses] one… then knows where the arbitrariness lies and

---

observer's perspective there is an ineluctable arbitrariness as to *which* universe they would end up in out of all manifested.

[9] Reprinted in Trimmer, 1980

> where improvement must be made in case of disagreement with experience: in the initial hypothesis or model… If in many various experiments the natural object behaves like the model, one is happy and thinks that the image fits the reality in essential features. If it fails to agree, under novel experiments or with refined measuring techniques, it is not said that one should *not* be happy. For basically this is the means of gradually bringing our picture, i.e., our thinking, closer to the realities.

Correspondingly, the consequence is that the state of knowledge is one of reaching equipoise or alignment with facts as arbitrary parameters are excluded and the truth is cornered into an ever narrower scope of possibilities. The possible narrows down into the plausible. "The classical method of the precise model has as principal goal keeping the unavoidable arbitrariness neatly isolated in the assumptions… for the historical process of adaptation to continuing experience." What is this "unavoidable arbitrariness", if not some limitations of the mind, the unfocused photograph, but rather the *deeper truth* which surrounds everything, the thing in itself which thought and measurement must accomodate and work around, the fog that is well-captured? Hence the question of the existence of a "blurry model" of physics, a reality which is inherently indeterminate and unknowable with definite exactitude. For after all some of the smartest people in history confronted this question and came up against this unavoidable arbitrariness and hit it like a wall.

Fittingly it would seem that it is itself uncertain and undecidable whether it is the camera or the object photographed which adds the blur. Critically, this undecidability would either imply a fundamental limitation of the human mind (cognitive closure) or the perception of deep reality in all its strangeness. What is unavoidable is what is most true. It impresses the conditions of possibility of knowledge onto the knower, what is come to be known, a fact, is an inescapable and indestructible "kernel" at the root of a phenomenon. Anticipating the notion of epistemological incompleteness, with regards to the classical (positivistic) model, Schrödinger says,

> Perhaps the method is based on the belief that somehow the initial

> state really determines uniquely the subsequent events, or that
> a complete model, agreeing with reality in complete exactness would permit
> predictive calculation of outcomes of all experiments with complete exactness.

Such dreams of total knowledge are not forthcoming, however, as "it is quite probable that the adaptation of thought to experience is an infinite process and that 'complete model' is a contradiction in terms, somewhat like a 'largest integer'." Though the phrasing is different, this prefigures the notion that there is always "more out there" and that the balancing of internal knowledge with external information is forever ongoing. Though the pursuit of its resolution is the goal of knowledge, ignorance is never perfectly eliminated. What all this would seem to suggest is that the world indeed best conforms to a blurry rather than a classical model.

To return to the card game example. Suppose that instead of guessing a letter on the back of the card, a player must guess a number from 1 through 10. They are given 3 cards. The cards are placed in ascending order. The player must guess the middle card by first flipping the cards on the right and left. Doing so allows them to adjust their expectation. If the smallest number is 4 and the largest number is 7, the answer must be 5 or 6. Each piece of information eliminates a degree of freedom, narrowing down the possibilities. An unlucky player may get a left card of 1 and a right card of 10, a hand that conveys the minimum of information. As a result of these flips the player is no more informed than if they simply guessed point blank. Alternatively a lucky play may get a left and right card with the difference between them of 2, meaning that there is only one choice for the middle card. The information on the back of the cards matter. Here the relationship between possibility and ignorance is highlighted. While one can be ignorant of both necessary and contingent facts, contingent facts—those that could go one way or another depending on some second fact—afford more possibilities for being ignorant.

Does this seemingly contradict the idea that it is the flip state of the card rather than what's on the back of them which is important? The fact of the matter is that it is impossible to not gain information by flipping the cards, and

uncertainty is therefore always at its maximum if the cards are unflipped. Regardless of the definite order of the series the middle card can never be 1 or 10. So flipping the left and right cards and getting 1 and 10 still conveys information, namely, it conveys that flipping the card flips were uninformative. Not all instances of perception have an information payoff. More pointedly, what we are concerned with here is not reasoning but perception. The digits on the flip-side of the card are not processed with raw sense perception but analyzed cognitively to infer the answer through deduction, whereas the act of flipping is purely perceptual. Of relevance here is how perception uncovers the inputs which are then converted into grist for reason's mill.

Flipping the cards is an act of perception. Perception does not presuppose what is perceived, the letter on the back of the card, or whether it will be informative or not. Nor does perception do anything with the information gleaned from it, instead it "passes it off" to cognition for processing. All perception detects an edge or boundary in a space against which the agent then constrains their movements. It structures what they do and where they do it. According to the theories of ecological psychology developed by J. J. Gibson, perception is *direct*. Rather than representing the environment, the organism's perception *specifies* it. The organism structures its relation to its environment by sensorimotor coupling, continuously adjusting its relations to it through dynamic coordination of body and world without the need for intermediating mental representations. (Richardson, 2008). The reasoning behind this idea of ecological perception is that there is no time to passively represent the environment, and the environment is never fully analyzed to such an extent that the agent would have a full enough picture of it for it to form a meaningful representation in time to act. Instead, there is fluid interaction between organism and environment as sensorimotor coupling recruits on the fly the salient, directly perceived aspects of the environment that are relevant to movement. While this notion is more relevant to sensorimotor activity and is not represented by the above card examples, the flipping of the cards is analogous to the direct perception of a structure which scaffolds the consequential outcomes of the organism's behavior.

Uncertainty is similarly directly perceived. The implication is a

counterintuitive reversal of the conventional idea that knowledge accurately depicts how the world is, that the goal of knowing is to slip into one-to-one correspondent alignment with the facts of the world. How the world is in its raw form is superpositional and maximally uncertain. Learning is the process by which all those alternatives are hewed down until the most frequent stable behaviors of the object become the definition of it. This knowledge is always brittle however, because any time we are surprised or see the known behave unexpectedly, it is merely showing another side of its innate potentiality for uncertainty. It proves to us that what we knew about it was merely only a surface gleaning of its underlying potentialities rather than the whole truth of the object. This view sheds new light on longstanding puzzles such as the problem of induction and inverts what our conventional understanding of what truth is. The truth is that the thing-in-itself is never perfectly grasped in its totality, that what we come to know is a statistically inferred amalgam of frequently observed behaviors which only speak to accessible surface properties of an intrinsically unknowable deep object. What is known is what is selected. knowledge is not the total representation of the object as a fully apprehended whole but rather our limited conceptual summarization of it . As Schrödinger (1980, sect 6) puts it:

> We are told that no distinction is to be made between the state of a natural object and what I know about it, or perhaps better, what I can know about it if I go to some trouble. Actually - so they say - there is intrinsically only awareness, observation, measurement. If through them I have procured at a given moment the best knowledge of the state of the physical object that is possibly attainable in accord with natural laws, then I can turn aside as meaningless any further questioning about the "actual state", inasmuch as I am convinced that no further observation can extend my knowledge of it - at least, not without an equivalent diminution in some other respect… We have nothing but our reckoning scheme, i.e., what is a *best possible* knowledge of the object.

Therefore it some strange but definite sense, knowledge is an act of distortion. To account for the mind's inability to process the true profusion of facts about the world and its objects in all their unaccountable interrelations, it must capture the most salient aspects of this hyper-complex stream as "headlines" and "bullet

points" rather than attempt to encode this intractable complexity into knowledge in a form that respects its true immensity. We then call those salient points the nature and truth of the phenomenon itself, disregarding whatever deeper hidden essence which out-scopes accessible knowledge. If this weren't the case, one would have to ask why what we call knowledge often comes up short. How can one be wrong about anything one claims to know? Why It is not because we are ignorant, it is because the world is ignorant; it is still making its mind up. And in the case of the collapse of the quantum wave function, it incidentally makes its mind up at the same time we do when observing the state of the particle. If knowledge with a capital K were so full and adequate, if the sum total of knowledge were complete, then there would be no more surprises. Some facts are however characterized by their *capacity to surprise*

Given that knowledge is always "facing down" the unknown, interminably adjusting itself to a seemingly inexhaustible reality, and given that knowledge is the result of thinking, how its it that knowledge comes about at all in a universe consisting of or cloaked in a veil of ignorance and uncertainty? How thinking relates to uncertainty in more definite terms will be the next topic.

The concept that there is a horizon to knowledge, which is delimited by an outer expanse of uncertainty that constitutes this kind of pervasive surrounding cloud or fuzz of doubt, is one of the main organizing principles to the discussion of ignorance and is an intuitive way to conceptualize ignorance (DeNicola, 2018). It's a unifying image that sets the stage for everything else. Here, we elaborate on this crucial idea.

To add some more flesh to these bones, consider Clark's (2016) comparison of thermodynamic models of intelligence to machine learning. "Thermodynamic free energy is a measure of energy available to do useful work. Transposed to the cognitive/informational domain it emerges as the difference between the way the world is represented (modeled) as being and the way it actually is" (p. 306). An agent's model of the world is assumed to be in a perpetual discrepancy with the world as is. Because the world is a thermodynamic system it is always changing and thus the agent's model is always permanently out of date. Through

a tightrope management of error-corrections and predictions, the agent balances its sensorimotor and cognitive activities to be in continual adjustment with the fresh information state of its world. Thus, "the better the engagements, the lower the information-theoretic free energy. (p. 305)"

Here the difference between cognition and perception in information-theoretic terms is highlighted. What is perceived is the free energy, what is cognized its its reduction. In other words, the agent confronts an unorganized, hazy cloud of objective uncertainty as it maps to thermodynamic change pin the phase space of its probable distributions. Beyond the results of its predictions is a horizon consisting solely of this undifferentiated mass of "informational free energy", an ensemble of undetermined bit values. By finding the right fit between perception and this phase space, the agent refines their model and reduces the subjective perception of chaos by bringing it under the control of a more accurate, updated representation. On this interpretation knowledge is a kind of shedding or exhaustion of excess possibilities. Such a regressive self-organizing process toward a stable state is implicated in deductive thinking and learning (Kirkaldy, 1965) . To act effectively, then, is for these two ends, perception and cognition, to meet. Ultimately, what one acts upon is not the raw sense data streaming in, according to the predictive processing model Clark advocates. Rather it is the built up predictive model that has accumulated and which is used as an efficient and immediate substitute for concurrent sensory inputs.

Consider the before and after sensation of playing chess for the first time without knowing the rules and then playing it after one has learned. In a manner of speaking, a definite boundary is crossed as these before and after states transition that resembles a phase transition, like water changing from vapor to liquid to a solid. The game looks patternless and inscrutable at first. Once the rules are learned, the game comes under representational control. The initial meaningless chaos tightens into an ordered framework and regularities come into focus. The function of each piece and relations between them sharpen and clarify. In a definite way this "tightening" of the board represents a loss of perceived randomness. All the ways the came *could* be played—with the queen moving like the king should and a knight like a rook—are replaced by how the

game is. This is the very form of understanding. Specific groupings of the board's pieces come together under the same rules, and the original, flat landscape of ambiguous, seemingly arbitrary elements fades out as a structured representation draws forward. Information-theoretic free energy, and along with it uncertainty, has been reduced. To the untrained eye, a complex object is an undifferentiated mass in which the parts are indistinguishable from the whole. knowledge kicks in when that equiprobable distribution of sensory states collapses into demarcated, discriminated components that can be seen as distinct in relation to the whole.

Plato viewed the process of obtaining wisdom as moving beyond particulars to universals, from uncollected sensory impressions to the universal schematic Forms underlying them. The idea bears some semblance to this crystallization of invariant representations or schemas from repetitive inputs based on a basic initial state of unorganized free information. knowledge is akin to *unforgetting* on this account, the emerging of a stable representation from *lethe*, the abyssal depths of unknowing as unactualized potential knowledge. knowledge is extracted from what is present, or as Heidegger noted, unconcealed, rather than asserted and attributed. knowledge is latent. To pass from ignorance to knowledge is to to activate a potential. What initially prevails before knowledge is established is more accurately *over-orderliness* than chaos, a lack of distinctiveness owed not to an absence of pattern but to too many patterns overlapping and canceling each other's meaning out. Over-orderliness captures the fact that chaos and order are not discontinuous but are really the same thing with varying degrees of lesser or greater processing salience. Pure randomness perhaps can be the only thing considered to be purely chaotic, but the vast majority of states under it express some degree of order. In the same way according to thermodynamics, all stable states are localized pockets of order juxtaposed to the prevailing climate of thermal equilibrium, but the two are always in a flowing exchange with each other rather than diametric and incompatible opposites. The limit of perceptible order is what is intelligible. To a mind equipped to perceive it seemingly random gibberish is semantic, and there is a connection between the flap of a butterfly's wings and a hurricane on the

other side of the world. To know is therefore to simplify, or put differently, to see no more than what is necessary. The association with knowledge and simplicity and ignorance and complexity will be a recurring theme.

According to this platonic conception, knowledge already exists embedded in an overarching and underlying plenum or manifold of a deeper, holographic, hidden reality in which all that is sensible is derived from certain universal, standard forms. Ignorance would then be a layer that separates the mind from these forms, like the regular stone of a mountain getting between a miner and its gems. The uncertainty enveloping and crowding outside the horizon of knowledge is an ontologically real and independent entity that supervenes on entropic physics. Epistemology, in this real and definite sense, is grounded in the physics and information science of entropy. knowledge is latent in ignorance in the same way that order is latent in chaos, in over-patterning. knowledge happens when things come together within a chaotic superset of patterns narrows to become an ordered subset as the agent gains mutual information relative to their environment .

It almost goes without saying that many if not all important decisions are made with some relevant pieces to the puzzle missing. That is why they are important—there are stakes. A problem is nontrivial not only because it matters to us for whatever reasons, but because the requirements for solving it are more demanding. There are no stakes if one has foreknowledge of the outcome, the choice is trivial. One simply selects the obvious. It's somewhat typical to think of decision making as "acting on what you know". On the contrary, the decisive input to the decision process is precisely what isn't known. There is then a certain statistical organization to the distribution of problems in terms of their triviality versus interestingness. The vast sea of trivial problems is "closer to equilibrium" of the most probable, while the interesting ones tend more towards the unique.

Decision making under suboptimal information conditions is determined in significant part by the nature of time and therefore the structural ignorance it imposes. We can only collect information that is available in occurrent time slices

moment by moment, and there is no a priori reason to believe these moments make available every relevant piece of information. Future information is noncommunicative with the present, and the past information is a fragmentary and diminished. Yet there is another dimension to this time pressure, considering we are constrained not only by what limited information made available in time, but by the time sensitivity of most problems. Most problems are characterized by their transitory urgency. Something must be done *now or else*. Like a college student rushing to complete a college exam before the clock runs out, we don't have all day to mull things over. Details which will inform the optimal decision might not even exist yet, because they await the development of the circumstances. Hindsight is revelatory in this regard. Everything makes more sense in hindsight, and this reveals an important epistemic principle: namely the *temporal contingency* of knowledge and, of course, ignorance. Much of what we are ignorant of concerns *when* the fact arises. This observation has a long history, notably the problem of future contingents going back to Aristotle. It suffices to mention it here because it is one of the taxonomic categories of ignorance, temporal ignorance, that is to be explained. This topic will be taken up again in its own right. What is of main concern now is how it relates to the central topic of this section, decision making.

Predictions, like any non-deductive inference, are fallible. One can predict, but predictions are always probabilistic in character, thus a lower grade of knowledge than certain knowledge. Information signals fly along with the unidirectional arrow of time that moves from past to future, so it is possible to discover things about the past and present but impossible for the future. So the future is "walled off" from us, and constitutes a huge part of this cloud of ignorance enclosing the mind, an unsurpassable boundary. We encounter life as a Heideggerian *thrownness*, tossed into the stream of events and struggling to discern as much sensible pattern in it as we can while forced to act as the world rushes at us in a way that sometimes feels like too much. If only we could know what our future selves knew in so many situations, life would make much more sense and most decisions would be trivial. One's behavior would be radically optimized and enhanced because the mistakes one will make become mistakes

one *would've* made had the foreknowledge of their causes were present. Of course knowing in the present what we will eventually come to know in the future is impossible outside of perhaps strange retro-causal quantum mechanical edge cases. Decision making under suboptimal conditions, or what Herbert Simon called *satisficing*, is a well studied phenomenon in the decision science literature that addresses the issue of reasoning and believing under uncertainty and incomplete information. Rationality is bounded, that is, restricted by the partially obscured nature of the problem, the information state and capabilities of the agent, and time pressures. The best decision is often not the most obvious, so a decision which satisfies and suffices for the bare minimum solution conditions will have to do. It is when we make do by refusing to "let the best be the enemy of the good." While research in this domain has largely concentrated on behavioral economics and operations research, it's clear that satisficing can be leveraged as an insightful epistemological concept more broadly.

In what conditions can we know everything we need to know to solve a problem up front? In any such scenario, the problem space is likely either trivial or deductive in nature. In the first case the problem is routine and simple. In the second case, it contains no surprises in the sense that like a chess board none of the pieces are hidden, it is an information-complete decision space. (The only thing you don't know is what move your opponent might make next.) The solution and the problem are both constructed out of the same syntax, to find the solution, merely rearrange the problem. The path to the solution is marked by the use of well-defined deterministic rules which lawfully always yield the intended result if correctly applied. To know the solution is to apply the appropriate given rules in the right order. In these artificially closed systems, global knowledge envelopes regional ignorance. This is an information complete problem space. At no time will the conditions contextualizing the problem change. The rules for math will not suddenly be revised, 1 x 1 will always equal 1. Such deductive information complete spaces may grow at their outer edges and be self-extended by future developments in the fields and novel recombinations of the principles already established. But nothing new that can come of it can contradict the old.

Any non-routine goal involving the open, real world involves working with

an incomplete knowledge-base and an understanding clouded by contingencies, what-ifs, risk, and unintended consequences. It is a situation where global ignorance contains local knowledge. Hence the well known problem of induction and the resulting iffiness of this important class of inferences. In deduction there is all signal and no noise, it is only one's own thoughts that must be worked out. One must just know hat and know how to apply the appropriate rules in the right sequence. In real life, the signal must be extracted from often considerable noise that is mixed up with the solution. The rules are not specified. Ignorance may prevail in both cases, in the case of some vexing unproven mathematical conjecture versus how to cure cancer, but ignorance in each of these cases operates differently and appears under different conditions. In the former case, it's generally understood that the solution resides in what is already given; some permutation of the infinite mathematical symbols must contain the answer. In the latter, the solution resides precisely in what is *not* given.

Contrary to the rationalist "mirror of nature" representationalist concept of mind, as Rorty (2009) labeled it, in which the reason has direct access to the world through pure correspondence relations and in which logic provides the means for a flawless reconstruction of natural objects in a conceptual language, the real picture is much grittier. This mirror of nature view is epitomized by logical positivist project, most conspicuously Wittgenstein's *Tractatus Logico-Philosophicu*s which posited a clear logical correspondence between ideas, language and the world and Carnap's *The Logical Structure of the World,* which envisioned a project of "rational reconstruction" in which a perfected language of science would allow for an ideal model of the objective world. The reality , argues Rorty, is not so pristine. Reality is under no obligation to conform to the framework of language. I would however depart from Rorty here and adopt a different take against the mirror of nature theory.

From the ignorance-first perspective, the world does not exist in itself as a complete, fully formed knowable structure. knowledge is not an mirrorlikek alignment with it, but instead a deviation from it, almost a mode of tampering, agent-centered use-making of information free energy. The world exists in the dark. To run with the metaphor, a knower is like a self-illuminating object, the

proverbial lightbulb of an idea, and what it comes to know derives from its own act of shining onto the darkened surface. Unlike a mirror, it does not catch the light. knowledge is never "acquired" in the sense that the outside world consists of a form that is commensurable with the form of knowledge and so pass into it. knowledge self-extends through the overcoming of ignorance. If it were to truly reflect a representation of the world, it would not be knowledge.

To make things even more difficult, we can only form concepts about which we are aware, and only some facets of the world are "submerged" as perceptible affordances accessible the senses. So incomplete internal knowledge must relate to incomplete external perceptual affordances, leaving exposed multiple cracks through which the fuller truth can slip by unnoticed. What correspondences can be made are then likely sparse and lopsided rather than an ideal one-to-one parallel mapping of a function to its codomain. Again there is no such salient co-domain of the world as a perfectly knowable object that has the same form as knowledge. Even the brain has functional limitations that bias and tamper with what it ultimately comes to know. The rational mind may indeed project correspondence relations to objective states of affairs, but it does not project a faultless logically pristine truth map onto the world. Instead it hoists its own contrivance on it which strives to conform external realities to its own agenda and requirements.

In decision theory these foibles are products of bounded rationality. One might ask what exactly bounds rationality, if not the sort of structural ignorance which informs the backdrop of all thinking, knowing and reasoning.

While some may think more clearly than others in many practical situations nearly everyone makes mistakes, and mistakes are proof of cognitive imperfection. The confusion stems from this fantasy that the faculties of reasoning are faultless and that the trouble with finding out the truth is restricted to the problem of how to relate the rational processes of reason to the facts of the world. Reason is a given, totally reliable conveyor of the truth and it is merely a matter of finding how to connect it with the facts. In the view favored here, reasoning is not concerned with the external world at all, but rather, with resolving thought. Reasoning is like turning down the temperature in a boiling

pot of water so that it settles and cools rather than overflows. Doubt and volatility must be diminished for reality to be perceived. Eventually some pragmatic fit or balance will be achieved, because the knower will be able to act and get what they want and by this basis they releive themselves of all doubt. The volatile struggle of competing hypotheses must eventually cancel each other out until only one survivor remains. The goal of reasoning is to put a stop to itself.

Everyday life is laden with examples of hindsight bias. Past decisions which were once seen as well-informed in retrospect now seem shortsighted. The answers which once eluded us now seem obvious. *How did I not think of that?* The best-laid plans ended up rudely rebuffed by an unexpected turns of events. As time goes on all these haphazard branches of wrongheaded belief and action get captured in a net of memory which reveals the truth by contrast. It is that net which we call experience, speaking in the sense of a veteran soldier, a seasoned chef, or a senior manager. *Expect the unexpected* goes the cliche, but this hints at a real truth that the unexpected is paradoxically the norm. Exceptions are unexceptional. The key problem here is that information gathering must stop at some point and a decision must be made. To continue to gather information is to defer the decision and risk "analysis paralysis." We can deliberate or we can decide; we cannot do both at the same time, and we cannot deliberate forever even though there is an inexhaustible possibility to do so.

This "cutting off" of deliberations ultimately exposes us to ignorance, a breaking off point. There may or may not be a definite formula for determining when the cut off should be committed to depending on the specific profile of the problem at hand. For many areas involving important decision making there is no clearly defined set of criteria the satisfaction of which can be taken as a signal to halt information gathering and commit to a decision. Most important problems are by definition nontrivial, meaning that the solution is not easy and obvious and accessible through a common routine. After all nothing tells us that we have *found everything* relevant to a problem until we act in a way that solves it, which is of little use beforehand. The premise presumes the conclusion. Unlike installing a software on a computer there is no progress bar telling us how

done we are, no external reference point to reconcile our internal information state with an external information state to establish the parity between them. Ultimately only an intuitive level of satisfaction is reached, the deed committed, and the result evaluated all on the basis of potentially inadequate information. Only when the results come in do we know for sure that our information was sufficient and the right decision chosen, which is small comfort when the result is negative and you could have made use of that information before making the decision. We are either proven right that we knew enough or the critical missing piece reveals itself and we exclaim "how didn't I see that?"

In this decision dilemma is a faint hint of the halting problem of theoretical computer science: how to determine when there is enough information to stop gathering more information? Like the halting problem, the matter is *undecidable*. We can't know with perfect certainty that we know everything there is to know about a given subject just as from an arbitrarily chosen program it is not possible to tell if it will halt or continue forever. We can only be shown that it we didn't know by being surprised by new information.

Only by structuring an investigation with a set of well-informed satisfaction criteria can uncertainty be semi-reliably abated and our thinking *uncertainty-proofed*. It is easier to eliminate sources of uncertainty external to the fact rather than to remove uncertainty to the fact: this is the principle informing the use of controlled experiments in science. Variables are carefully selected to minimize confounding factors. We have to contrive a method to come up with a set of mental checkboxes to check off given certain observational evidence. But this presumes a thorough understanding of the scope of the problem. Uncertainty can never be eliminated entirely. Aggravating the situation is the fact that most important decisions are *pressured*, either by time, resources, or some other dependency, which ensures that a decision must be made regardless if those self-imposed criteria are satisfied.

When it all comes down to it what isn't known will prove the most crucial factor in decision making. The emphasis on how the decision agent executes on the positive information content it has, removes the fuzziness from the equation, but all of this is a response to that fuzziness. Satisficing helps explain how the

best possible course of action is usually not explicitly prescribed by the current available information state. Understanding *why* satisficing is necessary entails an understanding the structure of uncertainty relating the agent to its information environment. There are dual limitations, external and internal to the agent, that impose uncertainty on it. Now to turn toward a more technical analysis of reasoning under uncertainty which implies some mechanism of belief revision.

In 2020 Facebook Research, a major player in the AI R&D space, published a paper on their new AI model called Recursive Belief-Based Learning (ReBel). ReBel was designed to learn imperfect information games such as poker, in which estimating hidden information is an important factor in performance. Conventional AI models have been trained around perfect information games such as chess and have come to perform at the level or better than human players. In games such as chess, all the possible moves and all possible states of the board and its pieces are afforded. Decisions operate with ideal salience, because every possible next state of the board can be derived from its current state with no hidden variables. Yet such programs struggle when relevant features of the information state are unable to serve as input because they are withheld from the current game state. For games like poker, and by analogy most of the interesting problems in life, relevant information is withheld or inaccessible and these imperfections must be worked around using guesswork.

ReBel weighs actions based on the probability that other actions have occurred. It is based on this idea of public belief states, which is the pool of information that has become available to all and which therefore constrains the scope of the hidden information. The model advances beyond discrete representations to *belief representations*. Discrete representations are mere snapshots of the current state of play and are the tip of the history of the game as it unfolded. Belief representation instead take that public information and use it to form representational structures that map the various probability outlays entailed by the available information. To demonstrate by example, suppose an opponent plays an ace, then that diminishes the possibility that they will play an ace again by 1/4. And if oneself has an ace, that probability is diminished to 1/2.

The program is not simply representing how the its current hand stacks up against an ace but also takes into account the remaining probability that an ace will come up again.) As the researchers state, "this illustrates a critical challenge of imperfect-information games: unlike perfect-information games and single-agent settings, the value of an action may depend on the probability it is chosen. Thus, a state defined only by the sequence of actions and observations does not have a unique value." (Brown and Bakhtin *et al*, 2020, p.2) Put differently, what is essential is what the past sequence of actions discloses about the probability distribution of the remaining optimal choices and how this public information adjusts the probability estimation of which choices are optimal. A public belief state is like a repository of the shared information which is not privately withheld between the players that allows them to infer what probability distribution of the remaining unrealized infostates of the game are.

To illustrate the idea, the researchers describe a variation of a 52 card game in which rather than viewing their private cards, the game is mediated by a "referee." As the authors explain,

> The referee then samples an action on the player's behalf from the announced probability distribution for the player's true private card. When this game starts, each player's belief distribution about their private card is uniform random. However, after each action by the referee, players can update their belief distribution about which card they are holding via Bayes' Rule. Likewise, players can update their belief distribution about the opponent's private card through the same operation. Thus, the probability that each player is holding each private card is common knowledge among all players at all times in this game. (Ibid, pg 4)

The key point is that the traditional card game which involves each player privately withholding their cards and iteratively disclosing their hand through a discrete action each turn can be represented equivalently using public belief states in this way. ReBeL is able to achieve the feat of representing an imperfect information game as a perfect information game.

While this research into imperfect-information games is provocative, but to

generalize philosophically from it caution must be taken. Real world epistemic agents face a universe of unknowns with which they are not engaged in a symmetrical interaction analogous to a card game in which both players are equally subject to the same context. The knower and the world of the unknown do not play by the same rules. Knower and the unknown are not dueling adversaries, at least in most cases. Rather, the knower is an intentional agent while the field of the unknown is a nebulous backdrop without agentive intent, though it may be interspersed with such agents of which it is unknown. The authors project that this research could be applicable to many domains such as negotiations, cybersecurity, auctions, and other multi-agent zero sum interactions. These are all adversarial engagements in which some parity is assumed between the driving components of the system. Nevertheless, as epistemic agents facing the world, we are indeed confronted by an information-incomplete situation. Except the conditions in play between the two cases differ drastically. Perhaps the fundamental difference between the isolated systems of information-incomplete games examined in this research and the general agent-world relation is that not only do we not know what hand the other "player" (the world at large) is holding, we do not know the complete rules of the game. Thus different "moves" of the other player strike us as at time anomalous and superstitious, lacking any context to make sense of them otherwise.

On the contrary, metaphorically speaking we carve out our own games, our own intentional, organized, constructed settings in which segments of the world are partitioned and subjected to rules, technologies, and practices which create manageable chunks of order. These partitions, the most obvious example being civilization itself as it is imposed over and above wild nature, come with their own uncertainties, but they are uncertainties which fall within brackets specifically chosen to facilitate rule-based, game-like engagements between people rather than arbitrary exposure to wide-open cosmic volatility of humans versus untamed nature.

In the attempt to flesh out this exploration of the varied interactions between

ignorance and knowledge many sources and perspectives can be investigated. The problem is multifaceted and there is little precedent for describing it. To that end Bayesian networks are a powerful model for analyzing the belief revision process that is prompted by a adjustment of mental models to a changing probabilistic world bracketed by uncertainties and the continual flow of new information. A Bayesian net is a probabilistic graphical model that encodes the conditional dependencies between a set of variables. As a directed graph each node in the collection has links to the others and has a relational structure such that changes in the weight of one node propagates to others, revising the overall probability estimate of the sum of each prior, the joint probability distribution. Each descendent node's probability is conditional on the join probability of its associates.

The basic work cycle of Bayesian net involves parameter learning. When unobserved evidence variables are introduced to the graph, an initial posterior probability inference is recalculated. Hence the net implements Bayes' theorem. Given the new evidence variable, the probability distribution updates. Importantly, the dependencies between the nodes inform the network and because of the networked nature of the graph it is able to inferentially "fill in the blanks" when data entries are missing (Heckerman, 1996) based on its recomputing of the probability weights given new input parameters.

The Bayesian network model captures deep truths about reasoning in a probabilistic, uncertain world. Each relationship between nodes, represents dependencies between beliefs, such that new information recomputes the degree of conditional independence between each assumption. Most notable is how the system encodes workarounds for data that is missing. In adjusts to its own ignorance.

In Bayesian terminology, probability is regarded as a degree of belief. Contrary to the classical conception of probability in which the chances of an event occurring is pictured as an intrinsic property of the relationships between the conditions themselves, Bayesian probability represents intentionality. It represents the level of confidence, the degree of belief, an agent adopts in estimating the outcome of events. As such Bayesian inference is a measure of

ignorance. It attempts to reconcile *how* much uncertainty remains after quantifying how each distinct belief about the event support each other taken together. Bayesian hypothesis formation importantly produces an estimation about what the system doesn't know. It's a generative model.

Bayesian nets present a fundamental theory about agent-world relations that says, "given what I believe about known conditions I can infer what is likely but that I do not know directly." Thus given the argument that unknown entities have just as much direct and active sway on states of beliefs as positive information contents, Bayesian networks are a good model for how the known grows into the unknown. The development of a Bayesian net is like the growth of a bubble. Beliefs are actualized over the phase space of all possible beliefs based on the sedimentary state of all previous beliefs summed together by their joint conditional probabilities. For minds afflicted by structural uncertainty as an inescapable reality of being in the world, such a model is an excellent and rigorous formalization of the belief formation process.

The flip side of the Bayesian tuning process is the downgrading of prior beliefs on the basis of new information that contradicts it. Parts of the Bayesian belief network can become deemphasized. New information can negatively down-regulate (to use a term from neurophysiology) the weights of the graph until those with previously high confidence are shifted to the periphery. This phenomenon poses an interesting question about the psychology of belief. If a belief that was formerly held in high probability is downgraded, it means that it must've never been true in the first place. It was subjected to doubt If that's the case, it proves that the belief system was never about what is true about the world in correspondence, but what was true relative to the set of beliefs themselves, coherently. This demonstrates how the coherence theory of truth—consistency as truth—can actually be a source of positive ignorance, While Bayesian nets can be used as models of strictly empirical analysis, they serve as a model of irrationality as well.

A relevant observation from cognitive psychology is confirmation bias and the resistance to new information that strongly conflicts with previously held beliefs. Where Bayesian models fall short is in modeling the specifically

emotional attachment to beliefs that people exhibit. The models are strictly rationalistic, they unbiasedly readjust their weights in the face of new evidence variables without the complicated motivational conflicts and biases that human probability estimation is subject to in the course of its conduct. In human belief revision, confirmation bias points to the underlying value-driven nature of thinking. Logically flawed human belief systems tend to allow new variables to influence estimates only if they are motivationally or narratively consistent with the previously held beliefs rather than simply logically consistent. Evidence variables must be permitted entry and evaluation rather than be automatically enrolled into the calculation. Scientific training might mitigate this tendency to favor consistency and reject deviation from priors which is so different from Bayesian models, but as we will see later in the study of cognitive bias, no one is immune from such fallibilities. This tendency usually has to be continually fought. "A wise man proportions his belief to the evidence" says Hume, but we are not always so wise.

When someone has the aspersion "know-it-all" cast on them, what is insinuated is that they are a really a know-nothing or a know-little. Their confidence underestimates the depth of the facts. They are gratuitously self-satisfied in their own limited understanding, lacking in intellectual humility, which enables one to continuously probe the rich depths of the unknown rather than to mistake current knowledge for complete knowledge. Intelligence is curtailed substantially once it fails to consider the unknown; to really be intelligent means to situate one's finite knowledge relative to what one inevitably does not know. It's not how much you know, but how you relate your knowledge to your ignorance in order to measure the latter, which defines truly effective intelligence. This discussion leads naturally to a discussion of a hypothetical cosmic know-it-all, Laplace's Demon, and how this very concept results in absurdity.

Various theological traditions envision an omniscient divine being. Far from idle, such contemplations seek to imagine what extends beyond human

understanding and envision what might be the very limit of intelligence. This a philosophically valuable exercise for epistemology, because it forces one to think about an intelligence far greater than our own, and as such posits an extremum at which the capacity for knowledge is absolute . Trying to comprehend the mind of God, conceived as an all knowing fashioner of the cosmos, is interesting because how starkly it contrasts with our own finitude. Epistemically, what is a god but a being which trumps all ignorance? What would it mean to be omniscient, to experience the world completely in the absence of ignorance? These questions are fruitful to imagine for a being which is decidedly closer to opposite end of the spectrum of knowledge capacity.

No one can even fathom what it would be like to know absolutely everything. So conditioned are we by finite knowledge, so used to confronting our ignorance on a regular basis, that imagining what it would be like to know everything is incomprehensible. One might imagine such mind-blowing knowledge would be almost incapacitating, as if some implicit recognition that we don't know some things is a comfortably familiar though flawed state of affairs. Or perhaps it would be defeating . To use pop culture reference about the stultifying effects of omniscience, consider the character Dr. Manhattan from the acclaimed graphic novel Watchmen. Among Dr. Manhattan's many superpowers is prefect insight into everything that has happened, is happening, and will happen. Such insight proves debilitating. He no longer sees the point of any course of action, already knowing the outcome well in advance. Lack of uncertainty defeats and preempts any meaningful reason to engage with the world. So Dr. Manhattan ends up teleporting to a barren alien planet and passes the time simply sitting and tragically contemplating the futility of existence.

God as a being existing outside of time is a different matter. Because what it means to be a being outside space and time is scarcely conceivable, theologically epistemic speculation ultimately lack a certain specificity and definiteness. The situation changes and gets a little closer to home when one wonders whether any finite being in the universe can gain perfect knowledge of it. As a matter of fact it is logical provable that *it is impossible for any part of the universe to completely know the whole in which it is contained*.

The great physicist Pierre-Simon Laplace brings to mind this conundrum with his idea of a theoretical entity that had total information of the state of every particle in the universe. Laplace thought differently. In a perfectly deterministic universe it is in principle possible to predict the future to an arbitrary degree of accuracy. He devised a thought experiment to illustrate his intuitions.

Conventionally dubbed Laplace's Demon, this entity has total and accurate information about every particle and its state (position and momentum) in the universe at a given time. According to Newton's classical mechanics, such perfect knowledge of deterministic states even of just one moment in time should provide enough context to predict every future state. Though the ensuing calculation would be vastly complicated, that snapshot tells you everything that will happen next, but if you know what comes next in a purely deterministic expansion, then the same initial information is sufficient to predict an arbitrary number of next steps out.

Putting aside questions of ontological indeterminism put forward by modern physics that rule out the possibility of deriving the exact future state of a particle system from an inclusive momentary snapshot of it at a previous time, the epistemological ramifications of the Laplace's thought experiment on its own is worthy of exploration. Laplace's Demon serves as a helpful mental model for understanding the true meaning of uncertainty and the limits of any possible knower. On this view, the incompleteness of knowledge is not just an expression of ignorance as a subjective state, but a basic consequence of being a finite being existing in the physical universe.

Mathematician David Wolpert's (2008) influential impossibility proof of an inference engine of Laplacian scope sheds light on the problem. Using a variant of Cantor diagonalization method, it shows how perfect knowledge of the world is impossible and thus that ignorance is inevitable. Even in strictly deterministic, classical universe, a perfect inference device existing within the same physical universe it observes is impossible. Wolpert posits a "non-quantum mechanical uncertainty principle" which is evocative of the ontological uncertainty we seek to characterize.

In a purely physicalist universe this complete knowledge of particle state

should extrapolate to a description of every fact in existence. Knowledge of every fact supervenes on the total physical state description of the universe. The demon would not just know where all the particles are going and will end up, it will know about everything that involves those particles, every fact writ large. If however, the demon is counted as part of the physical universe, it then falls prey to specific problems in set theory and Russell's Paradox. If such a universal set were to contain itself, this would merely create a new superset of size n+1, where n was the previous size of the set before the entire set was included as an element. The set of all sets must contain previous iterations of itself infinitely many times . Each embedded element of which is infinite, because it must contain each of its subsets, which contain every element of the set, including itself. Thus it can never arrive at a complete universal finite set. As the observer, the demon is necessarily part of the observed physical system. As such it does not execute an inference magically, it requires a mechanism to implement it, or what Wolpert calls set up function. Total self-knowledge requires a recursive dive into the demon's own model of the universe, or *meta-knowledg*e, and so on in an infinite loop. If the demon is thought of as a program, this infinite loop would break it and ensure that it can never answer certain questions involving its own knowledge state, as the halting problem dictates. Its model of the physical system is a separate entity unto itself whilst also comprising part of that physical system. If the demon is part of the physical universe itself rather than outside of it, its own knowledge state cannot be untangled from the rest of the physical state that both surrounds and suffuses it. This, as we shall see, is a variant of the liar's paradox, and it becomes clearer once two demon-like agents are posed against each other.

Wolpert's exposition is highly sophisticated and mathematical. Rather than reproduce this deeply involved proof, we will here instead extract from it its overall implications. First, several basic assumptions must be assumed of the inference device. The device must be counted as part of the physical system it observes. It must be configured, like a computer program or experimental setup. It must therefore instantiate setup functions which correspond to the variable it observes. And it must be capable of being posed yes or no questions. Such a

device is limited by the fact it is embedded in the system it attempts to predict. The device runs up against *inference complexity* that approaches infinity. In essence, the collection of set up functions the device needs to initialize to approach absolute predictive power becomes prohibitive. Such a device ought to be able to answer any yes or no question asked it, but according to Wolpert there must exist at least *one question* that it is impossible for it to answer.

What is this question and what are its properties? There is an aspect of self-reference. Similar to the liar's paradox, the system cannot answer a question about what its answer to another question will be. A question such as "will your answer be yes?" is unanswerable, because to do so would require that it initialize a setup function for another setup function that points to the observable, ad infinitum. a prediction concerning the narrator's future cannot account for the effect of the narrator's learning that prediction," Wolpert claims (quoted in Marshall, 2018). The device gets bogged down in its own complexity and inference generating machinery.

Here's the kicker. If more than one such device exists, such a device is impossible. It's ability to read into the other device's workings becomes entangled in the inference complexity of that device. Put differently, "a reality can have at most one of its devices that has free will and can predict / recall / observe the other devices in that reality with guaranteed complete accuracy", says Wolpert (p .23). If you pose the same question to both devices, neither device is able to predict what the other device's answer will be as well as its own answer to the question. Assuming super-turing computational powers, the other inference device acts as a kind of inscrutable black box to the other, its very powers of inference making it impervious to prediction.

Even total knowledge of every brute fact entails a new set of unknowns because there must be a metacognitive ordering and registration of those facts. Laplace's demon is a mere "inference engine" it is not a mind that dwells and reflects on its thoughts in the conventional sense. If the demon doesn't know how it knows, as Wolpert's theorems imply, it can't decide what to do with all the information in the world because it is swamped by its requirements for gathering information about itself as a part of that same world. In order for it to predict

there must be something it doesn't know. It must work through an infinite amount of steps to get there, or take longer than the timespan of the universe to compute.

Laplace's Demon's disproof is instructive as a limit to absolute knowledge. Its logical impossibility establishes the the claim that there is in fact beyond doubt a category of absolute ignorance that precedes knowledge with mathematical precision. It's not uncommon in philosophy and in thinking in general to presume the subject-object distinction. This partitioning of the universe into two fundamentally incommensurable divisions, the observer and the observed, reappears in many guises, from the mind-body problem to Kant's phenomenon/noumenon. The universe however is one thing, and the knower is part of it, and thus its capacity for knowledge is ultimately constricted by this fact. Wolpert's proof demonstrates how these considerations place restrictions on what any finite physical being can come to know about the world that envelops and constitutes it. A being outside spacetime, or perhaps occupying a higher physical dimension as in Edwin Abbott Abbot's tale *Flatland* might be excused from these physical constraints. Such a being is purely hypothetical, however, and the main focus of this study is centered around an empirical framework. Later in Part II the deep significance of the conditioning of knowledge by physics and therefore inevitably also ignorance, because knowledge presupposes it and vice versa.

Review so far

It may be helpful to take a step back and to take stock of what's been established so far. The major idea is the knowledge is conditioned by ignorance. Ignorance is fundamental to an epistemological universe, one that is essentially information-based. knowledge and ignorance are not distinct from each other, oil and water, but rather, the head and tail of the same ourobosis, or continuous. The positive content of assertive propositional knowledge is shadowed by an inverted, negative space of ignorance, which it "struggles with" in doubt and which it is "driven towards" by learning.

Ignorance isn't simply an absence of knowledge but has a detailed structure and taxonomy. Lining the boundary between knowledge and all the possibilities that reside beyond it is a buffer of uncertainty. Empirical methods are the best way to probe past this buffer and scrutinize new potential candidate knowledge items. There are other methods that will be explored later. The flaw of epistemology historically has been its over-concentration on the issues of the positive content of knowledge, what knowledge is, how we know, and how to define knowledge, but less about this buffer of uncertainty and what lies beyond it in relation to ignorance.

Particularly when we visualize knowledge acquisition as analogous to a thermodynamic work cycle, the pull of ignorance and the conscious desire to eliminate it which drives directed learning, the relationship between knowledge and ignorance sums up crisply. knowledge acquisition is essentially goal directed behavior that intentionally objectifies ignorance and seeks to eliminate as many of these undefined variables as it can identify by excluding excess possibilities. The process can be compared to completing a jigsaw puzzle, and deducing what pieces go in the empty spots by comparing the hole to the surrounding assembled puzzle and the remaining pieces.

To elaborate on last point about minimizing variables consider developmental psychology experiments. To the child's point of view the way an experimental set up, say the famed marshmallow test, is set up has no bearing on their understanding. They simply see what is presented at face value. The series of experimental variables are taken as if not uncaused, then not related to what they are experiencing. There is no background to this naive realism. In a similar way many narratives or surface explanations (folk psychological theories) can float freely and independently from true causes, the hidden variables at the heart of the matter. This may not stop us from operating in our world however, because if others believe the same things, a certain cognitive economy forms around these fabrications. In a manner of speaking knowledge representations project a *distorted by workable* rendition of the world suitable for our purposes, just as a bee sees ultraviolet light radiating from a flower in order to guide it to its objective. One could argue that the bee does not falsify the object of its senses, but rather

perceives a cross-section of the underlying reality. However, it would be a mistake for the bee to believe what it sensed was a complete summary of the object.

It's gives one a feeling of deep insecurity to recognize that beliefs can be so "untracked" from the actual state of affairs but assuredly much of experience is framed by such cognitive manipulations. For example, folk psychology explanations for other people's behavior can have valid heuristic value but be completely divorced from the facts that do not subtend the boundary that it has imposed and which all it knows falls under. Consider architects in ancient times who constructed magnificent buildings without a precise understanding of the underlying laws of physics that formed the Kantian *conditions of possibility* for such a structure. Or consider the once inconceivable idea that a chunk of uranium could be harnessed to release unfathomable nuclear power that is somehow mysteriously packed into such vanishingly small atoms. Nobody likes to acknowledge that their intuitions are a front and that there are hidden causes behind the curtain pulling the levers. But the pursuit to resolve this discrepancy between appearance and reality is the motivation of science. Science is about bringing intuitive, flawed, all-too-human (mis)understandings into reconciliation with the weird, counterintuitive, inhuman reality of nature. Science is a method of *learning by at first unlearning*, of overcoming deceptive intuitions. You can be sure you are making progress if your bubble bursts.

To "burst one's bubble" is a particularly evocative phrase with respect to this epistemic theory. Having a mental bubble in the first place is to have a consistent set of beliefs, which derive their structural integrity from their relation to one another rather than to the facts. To have one's bubble burst means to encounter facts which are so overwhelmingly evident that the network of beliefs cannot stand before it. And so we learn by unlearning; the previous belief structure collapses and gives way to a more reality-approximate estimation. So long as bubbles keep bursting learning is happening and there is health and activity to a mind. Once no more response to surprisal (information value) is experienced, the system has effectively senesced. It is stuck in its ways and incapable of further growth.

If complete knowledge of the world is impossible, knowledge acquisition is inherently a *negentropic* —entropy reducing-process. To proceed at all in the action of its mechanism a cognitive engine must operate on a substance of unknowns. Unknowability corresponds subjectively to a random, disordered, high entropy state and knowability to a predictable, ordered, low entropy state. That very disparity between known and unknown is similar to to a gradient in physics. In an equilibrial state there is nowhere for energy to diffuse, and so all activity in the system evens out and slows to a minimum. Investor Warren Buffet had a point when he quipped, "if you're the smartest guy in the room, find another room." If there is nothing more for you to know in a situation, if you've learned everything there is to learn, you will eventually stop thinking. Without challenge there is little progress. You have reached a kind of stagnant epistemic equilibrium.

knowledge is represented physically. Another person can put knowledge into the world that isn't yours. Through the establishment of educational traditions an environment is structured so that public knowledge can be placed in a pool where it is made available to trip student's ignorance/knowledge transduction mechanism, colloquially called learning. In such epistemic environments, ignorance is thinned to a minimum by supplying students with the right things to learn that they are initially ignorant about. Once again, the acknowledgement that the student is ignorant of the curriculum forms the rationale for them to study it and is the source of propulsion motivating them to succeed in the course.

A consequence of this line of reasoning that will be expounded on in greater depth elsewhere is that the connotations of ignorance as always being a bad thing are not absolutely valid. In certain situations it can be good to say there is a lot of ignorance here, in that measuring what you don't know in relation to what you do is the first step in learning it. You have to be aware that there is something new to learn to learn it. The truly bad kind of ignorance is of the second order, ignorance of ignorance, because this type of ignorance is not actionable. It's pure obliviousness. In such circumstances one is apt to cling too tightly to what one already knows, and to mistake it for "all that is in heaven and earth."

It is a simple and uncontroversial fact that there is more that is unknown than

known. This is, always will be, and always has been the case. I do not know the position of various asteroids hurtling through the void of space lightyears away, of which there are surely millions. Neither do I know the most of what is happening on this very planet, or even my own neighborhood. I cannot fathom all the particles in a single breath or what is going on in a single cell on my thumb. And despite the expanded information capabilities afforded by the technology, it offers only a fractional sample of the ambient state of global affairs. Neither do I even know how my own mind functions to a degree of perfect confidence. This greater set of unknowns includes all that was, is, and will be, that falls outside my particular learning path history. The unknown then becomes the backdrop, an actual substrate and raw element, which precedes the known. It still exerts a pull on the known, like undetectable dark matter exerting an appreciable influence on observable matter.

This enters a new phase in this discussion about the interactions between knower and unknown. The main argument here is that the central subject of epistemology should be this relationship. These two categories are in dynamic tension which the epistemic agent seeks to reconcile through learning in discovery when knowledge is in the ascendant, or who falls prey and is benighted when ignorance gains the advantage. This full picture cannot be appreciated however if ignorance is exclusively treated as a simple lack of awareness or the absence of knowledge rather than a rich substance with a detailed structure in relation to the (un)knower.

The Interaction between knowledge and Ignorance

Nothing in epistemology makes sense without both the relation of a knower to knowledge and in relation to ignorance. Knowledge does not reside in a perfect field of pure knowledge, but always in reference to a epistemic agent typically with incomplete access to information and always with manifest knowledge . That agent by extension always exists in relation to ignorance, which is what conditions and limits their knowledge. The agent does not have to

be aware of their relations to the world for those relations to still hold. Hence, what we are ignorant of still plays an active role in shaping our experience and thinking. This knowledge-ignorance relation can be described in various ways, either using neural or computational models, or cognitive-psychological theories and philosophical analysis, and other tools. In any case the underlying logic is the same. When at our wit's end, when confronted with something we simply can't wrap our heads around, we hit what feels like a brick wall. Acute and sometimes conflicting sensations of frustration, confusion, and curiosity afflict us. This protean, uncollected state of self-aware ignorance is a moment of dynamism of paramount interest. If we ever break out of this deadlock, either painstakingly or through sudden flashes of insight, new knowledge finally emerges from the unknown or put another way potential knowledge becomes actual.

That state of ambiguity and striving to overcome ignorance is the moment of victory or defeat before the unknown. It is the scenario in which the implicit Bayesian processing of the uncertain is conducted. Only when confronted by the limits of our knowledge is when we have a chance to expand knowledge. However not all beliefs are useful or correct, and what beliefs and knowledge states are accessible to us at any given time is in large part path-determined by what we have already believed and learned up to that point. In an odd way knowledge can be the cause of ignorance if our past knowledge constrains what we can know, fixing an ever more narrower area of comprehensibility that nevertheless grows more salient in time. There is an *opportunity cost* to learning, as you have no choice but to exclude other things to learn. Even the acquisition of knowledge therefore implies a corresponding gain in ignorance elsewhere, because one could have spent the same time learning something else, and what one can ultimately come to know is a function of how much time one has left to learn more. This point is illuminating if we regard how it can drastically bias how a person comes to understand the world. This is another counterintuitive example of how the possession of certain knowledge can lead to another type of ignorance. There is even a term for it in the literature, the "curse of knowledge" for when experts falsely assume that non-experts have the same background and

context. Knowledge can be overly constrictive, causing one to see the world lopsidedly. "When all you have is a hammer everything looks like a nail." Dennett (1995) calls such an inclination greedy reductionism. The thrust of the reductionistic principle isn't necessarily wrong on its own terms. It goes wrong when overly and hence greedily applied beyond its proper scope.

Cognition is a duel with ambiguity. As Clark (2016) writes, "neural representations… encode 'probability density functions'… and the flow of inference respects Bayesian principles…It means the nervous system is fundamentally adapted to deal with uncertainty, noise, ambiguity." The brain is thus viewed as a predictive engine that generates a theory about the world based on the degree to which sensory feedback data from the sense organs match feedforward expectation signals. Meanwhile discrepancies between feedback and feedforward signals triggers a work cycle in which the cortex updates its predictive model to accomodate the surprisal value of unexpected input.

Given that epistemology equals a theory of knowledge plus a theory of inquiry plus a theory of ignorance, you cannot have a theory of epistemology without all three factors and have grounds to make categorical, and ironically self-contradictorily absolute, existential claims such as all knowledge is relative. Neither for that matter can you understand the true meaning of these terms without reference to each other. Claiming that knowledge it relativistic is just a misleading way of saying that it can never escape the shadow of ignorance. Lastly, there is nothing relative about the specific knowledge that knowledge at large and in general is incomplete, it's a brute fact as objectively and universally true as anything that must conform to the rules of logic.

The question remains: how does the knowledge acquisition process and the general thinking process relate to an agent's limited perspective of the total world? No matter how broad one's personal experience or learnings, it is an irreducible fact that one only becomes acquainted with a vanishingly small fraction of the whole universe. There is always more in heaven and earth. The problem of induction therefore beckons. How do we so confidently generalize over a universe of possibilities that we have such a fragmentary acquaintance with?

The answer I would like to suggest, is the counterintuitive affirmative, active and real status of negativities, specifically those of the uncertain and unknown in their influence on our thought and activities.

Physical entropy and the subjective feeling of mystification share an exact homology. Both bring to mind a random, confusing anti-pattern, a fog of chaos about which the only certainty that can be said of it is that it is unpredictable. Thermodynamics regards the behavior of gasses as intractable at local scales but as obeying laws of general tendency when ensembles of very many particles are taken into consideration. The only inferences we can make about gasses are about those that involve very many particles. The local interactions between small groups of particles is inherently chaotic because they "bounce" off each other willy-nilly. By definition a mystery is no mystery if you know of an event just like it. For any entropic microstate, there is a high probability or perhaps even certitude that it is the only state like it ever to occur in the history of the universe. Although theoretically it is possible for the micro states to arrange themselves in an almost crystalline order, this is vastly improbable. The thought experiment called a Boltzmann brain was devised to think about the mathematical possibility of a set of diffuse particles spontaneously organizing themselves into a fully realized human brain floating in outer space.

Order in contrast is predictable. meaning that a repetitious temporal stability in structure can be identified. knowledge, like particular stable physical systems, is said to *crystalize*. In the same way experience of the unknown is a highly entropic state. A situation or problem is jumbled beyond recognition with no predictions available to crystalize the episode into an ordered pattern, just like in the chess game analogy. Ignorance is the analog of entropic randomness while knowledge low entropy order. Indeed, the hallmark of knowledge is often, and especially in science, identified as predictive power, whereas when the answer is unknown, when the issue is undecided, it is "up in the air", hanging in a state of uncertain suspension.

Thus begins a process of reducing that as of yet uninterpreted, structureless, entropic unknown and framing an understanding of it so that knowledge is crystalized. knowledge acquisition then reveals itself as a negative, subtractive

act rather than a positive productive act. To produce knowledge you would have to construct it from knowledge, which assumes regressively that you were already in possession of it. The unknown is negated, and by that negation, knowledge is left over. The unknown, conceived of as hanging over and obscuring uncrystallized possible facts, is *chipped away at* by the analytical procedures of investigation until the fact which it hides is "left over" like the remainder of arithmetical division. The mind is like a sieve that separates the dirt from the gold nuggets. There is therefore a negative character to knowledge acquisition. Because the unknown is a real quantity, harkening back to Wolpert's "non-quantum uncertainty principle" and to the negative definition of information, to eliminate it is an inversely productive act. The reduction of uncertainty is the negation of a negation, which yields the eventual positive content of knowledge as a sort of evaporative condensate, byproduct or remainder. More will be said about eliminative cognition further on.

Ignorance is a direct information-theoretic mapping of physical entropy. It's revealing that information theory makes such ample use of entropy as a concept. Entropy corresponds to the noise of the signal. Information loss is associated with the diffusion of energy or interference in a transfer medium. Shannon's definition of information is telling. "Reduction" means to take a way, make less of. How then can one reduce a nonexistent *un*-certainty? Uncertainty should not be a thing that can be operated on if it is merely a mental abstraction. Intuitively you would think that the definition of information should be that which produces certainty but the world doesn't work that way. Certainty is elusive, and we never know that we have acquired all the information necessary to produce an absolutely certain conclusion. The world has a probabilistic, open-ended nature both on an objective quantum mechanical and a subjective experiential level. What we experience is a series of anticipations of varying accuracy that could be rebuked by new information at any instant. Hume's saying that it is by "custom alone that we think the future conformable to the past" is another way of saying that prior predictive information never amounts to certainty no matter how much of it is collected.

It is therefore more correct to say that uncertainty is reduced rather than

certainty is produced, because perfect certainty in inductive inference is impossible.. Even such obvious commitments to certainty such as that one is not presently on the moon or that every person has two parents is bordered by a fine line of uncertainty. You could hallucinate that you were, you could be suffering from a serious psychiatric delusion, on mind altering substances, or having a dream. You could be a brain in a vat. Such uncertainty may be utterly, vanishingly minuscule, such that it barely registers as a doubt. But as, Wittgenstein in *On Certainty* asserts, one can nearly will oneself to doubt it. As one moves closer to less obvious statements of certitude, that faint, peripheral static of uncertain noise grows louder and more pervasive. Suddenly more of the fact is obscured and distorted by alternative possibilities demanding the assignment of a probability distribution to declutter and segment it. Most knowledge, being perishable and subject to revision, therefore exists in a state of fragility, always susceptible to a sudden spike in the unexpected that perturbs and transforms it. We are even ignorant that our knowledge might be flawed or has an expiration date.

The result of these conclusions match the predictive encoding model of perception. Cognition is built on perception, and perception is a tool to cut through uncertainty. So if perception at its lowest and most granular neural processing traffics in probabilistic estimation, so should cognition at large. Hence Bayesian belief networks entrain external inputs. They constitute a set of mutually supporting inferences that act together as a relatively stable backdrop for dampening as much uncertainty as can be achieved.

What properties or conditions of the knower make certain things knowable to it and others unknowable? To consider how the knowable is conditioned by the knower is to bring up the question of cognitive closure or boundedness. Are there facts which cannot be known because of strict limitations on the human mind? This question will receive a proper in depth treatment later on, but it deserves mentioning here.

There are of course cases such as dyslexia, prosopagnosia, color blindness, myopia, and similar dysfunctions which impose structural ignorance on a

person's perceptual capabilities. These impairments distort an agent's information access. Consider our inability to hear certain sound frequencies or see ultraviolet light, or to envision with the naked eye the microbial underworld teeming everywhere around us. These structurally imposed defects of perception contrast with the more domain-general issue of intelligence limits, or to be blunt, *stupidity*. Consider that the distinct wiring of any individual's brain determines the way any incoming sensory signals are interpreted. Sensory information only flows up pathways that are available for it to go through, the predetermined neural structures that have been shaped by prior experience. If there is not the appropriate structure for it to flow through, the information won't be properly interpreted. The specific learning history-determined organization of any particular brain *personalizes* any given percept. Where one may look at a math formula and see gibberish another sees meaning. The truth of the matter is that both of them are right in their subjective take, because the symbols only make sense or not if the appropriate structure is instantiated in their brains to interpret it. The deciding factor for the meaning of the symbols is not bound to the object, the formula itself, but to the distinct wiring diagram of each brain and how it realizes the appraisal of the stimulus. This type of ignorance I call *relative ignorance* because it is not absolute but relative to a given agent's personal experience and limitations.

While science does not yet know how to decipher how one patterning of neural connectivity brings out one worldview while another a different one, we can still trace the causes of those neural differences by proxy by looking at each person's upbringing and eduction. It's a mistake, however, to say a person's upbringing and education makes them believe and act how they do. This is a proximate cause. The immediate cause is their unique neural connectome. Past experience determines future experience by selecting what parts of the world are most salient to the agent, leading to a cycle that is both expanding and contracting, additive and narrowing, compounding and filtering. Every choice is necessarily exclusive of other alternatives. *Choice entails some information loss, an element of ignorance.* We will *never know* what might have been. We can't even know if if there is such a thing as might have been, that is to say, that potential is

latent to the world or we are "railroaded" to a strictly deterministic fate. The more attuned and attentive to given perceptual features, the more the brain's processing and memory shapes around it and becomes "fitted" to that content, to the possible exclusion of others contents. This fitting contributes to what the knower can know in the future and what cannot be known. For one who cannot speak or read Chinese, a section of thoughts and nuance with distinctly Chinese bent are inaccessible to them. Time spent learning to play the piano is time spent not learning how to learn something else. Thus, what we know draws forth what is knowable to us, certain regions of ignorance become more probably knowable based on the bias of the knowledge base. Meanwhile, ignorance in other areas grows more entrenched.

Commonly it's intuitive to view the relationships between knowledge and ignorance as strictly adversarial. Ignorance blocks the further acquisition of knowledge. Such conflicting dynamics are only one type of interaction. As has already been pointed out, ignorance precedes knowledge and is its prerequisite. Ignorance can in some cases serve as potential knowledge rather than an impediment to it. We will now explore these nuanced relationships in turn.

An often overlooked important background condition of epistemology is the lack of a specifiable universal point of reference from which one can make categorical statements about the knower, how they know, and their knowledge. Sometimes unrealistically perfect or universal conditions for knowing are assumed. Often an idealized and imagined knower is implied, "the knower", but this the specifics of this agent and the environment in which they know is rarely contextualized. Beyond even that knowledge is often referred to as though it were some pure abstraction in a void existing independently of a substrate with definite parameters. "To know" is not an automatic, transparent and intuitive uptake of information or facts that operates the same for all possible knowers. A lot is going on under the hood, and what is going on varies case by case.

Common sense tells us that both the *amount* and *contents* of knowledge varies from person to person. Yet there is more homogeneity in *how* they know. If they know, they know in an identical way. I am certain that doctor knows more about

medicine than I do. Yet if both of us know that high blood pressure is bad, we share the same piece of knowledge. Though the doctor may have a more detailed understanding of why it's bad, that just means they have more knowledge. It's a difference of degree. But is there some shared strict upper limit on how much either of us could possibly know (according to Wolpert, the answer is yes)? Or does this capacity vary between individuals based on intelligence or some other variable? Or are certain knowers able to have a fidelity or clarity or depth of knowledge about a given fact that is somehow qualitatively superior to another? Very plausibly, seeing as on one level there is as much variance in knowledge between individuals as there is variance in experience between individuals, which is to say total variance because everyone has a strictly unique spatiotemporal experience of the world. There are also are various shared overlaps in knowledge between people, and correspondingly other meshes of relative knowledge and ignorance across populations. Is a fact that is well known by many "more known" than one that only a select few is privy of?

A common assumption of epistemology is that knowledge about a given object is the same in the mind of every knower. Anyone who knows x shares basically the same mental state or propositional attitude towards x. What is under-appreciated is that in different minds, a sort of amplitude or intensity of knowledge can be lower or higher, dimmer or brighter. Or perhaps they know in completely dissimilar, but equally valid ways. To those of great intelligence the same knowledge—the same reasoned belief in a fact—illuminates more and its usefulness is greatly expanded. Not all knowers do the same with knowledge. Mathematician John von Neumann, recognized by all who knew and worked with him to be a genius of the utmost caliber, had a mind which even to his lauded mathematical colleagues seemed to operate with otherworldly, superhuman capabilities. Critically, Neumann and his colleagues often knew the same things, it was rather their cognitive operating systems in which this knowledge was enclosed and situated which differed and for which von Neumann was blessed with superior resources. Of particular interest were two traits von Neumann was known for, the *speed* of his thinking and the efficiency of

his memory.[10] Processing speed and memory capacity the two benchmark computational resources, and generally the faster the processing and the greater the memory storage the better. Whether this marks a qualitative or quantitative difference or regardless if von Neumann was only distinguished by the fact that he was able to do what his colleagues did an accelerated rate, is not the point. The key issue here is that, even though von Neumann and his expert colleagues knew the same things, did they truly have the same knowledge? Or did von Neumann's gifts, his enhanced cognitive parameters which contextualized that knowledge, qualitatively transform it? Von Neumann seemed to exhibit a categorically different gift for mathematics, his knowledge of it went beyond a mere collection of learned methods and formulas but resembled more a kind of natural instinct or drive he possessed.

This goes to show that we can't always assume referring to "the knower" refers to something well-characterized from which valid generalizations can be inferred. Intelligence and knowledge are deeply related, but whether it is possible ground an adequate epistemology without respect to the intelligence that conditions it is not a given.

Another tendency is anthropocentrism when making categorical epistemic claims. Is there any parity between what your average human and a hypothetical super-intelligent AI might know? Is this difference a matter of degree or kind? How about a person and a sapient alien species with an entirely different cognitive architecture and perceptual apparatus? Must all sapient creatures be able to know at least one of the same things? Can a dog and a person know some of the same things? During a hunting trip, a hound dog perks up at the sight of a fowl in the bushes, while the hound master turns their gaze and attention toward it. Both discover the bird's presence. Is the dog any less knowledgeable of the presence of the bird in the bushes in this case? Is the human, with their impressive super-charged brain, doing anything substantially different? Perhaps both are merely *aware* of the bird, but the distinction between awareness and knowledge is not clear cut.

---

[10] https://www.cantorsparadise.com/the-unparalleled-genius-of-john-von-neumann-791bb9f42a2d

These variations in context are rarely appreciated in epistemic logic. The prevalent tendency is to presuppose idealized, perfectly rational conditions of airtight reasoning and information. That is not reality. In the pursuit of such a distilled, decontextualized and abstracted epistemic paradise ironically all that is significant for understanding the naturalistic conditions in which knowledge is implemented and applied is pushed out of sight. While I don't deny the value of epistemic logic for clarifying thinking about knowledge, there is something ironically ignorant about this reductive approach. Can any absolute claims about what is knowable be made without some assumptions that at their roots taper off into ignorance?

One thing is clear, and it points to yet another revealing asymmetry between knowledge and ignorance. While knowledge may vary in its forms and expressions and while knowers have different capacities, everyone is absolutely ignorant in the same way. Taking into consideration the pure uniformity of ignorance, anyone who doesn't know doesn't know in the exact same way as anybody else. They may only differ in terms of why they are and what they are ignorant about.

Frequently what is also unspecified is the background framework in which knowing occurs. A single instance of knowledge relies on a vast, undeclared, obscured infrastructure of epistemic dependencies. An entire network of implicit and often unacknowledged presuppositions are at play when it is asserted that such and such is known. The very foundations on which knowledge rests are submerged in ignorance because this framework of dependencies is often implicit. Only through technical rigor such as a logical proof theorem are all assumptions explicitly stated, though one might propose that even then there are hidden dependencies, such as that the formalism used is sufficient to prove its assertions exhaustively. We might say it is impossible to know what the first word ever spoken was, but that is certainly not unknowable in principle. It is only unknowable from the doubter's standpoint in time relative to the fact. This ascription of epistemic impossibility presupposes an arbitrary and implicit frame of reference for determining the conditions of knowability. It suggests that what is unknown must be bounded by what is already known, but if it were truly

unknown that couldn't be true.

By extension, there is no way to specify based on what is currently known what can be known or how much remains to be known with perfect certainty. The "state of the art" is just that, a state. Soon to be replaced by another. The cutting edge is cutting into something more and will presumably keep cutting. If it hits a snag, that could either mean that this pathway of discovery has been exhausted to completion or everybody has run out of ideas as to how to advance. How is it possible to know that there will never be a new development in a scientific field, that the world will be exhausted of its secrets? In *The End of Science* by science writer John Horgan, he claims that in general outline, all the fundamental theories of nature have been discovered and articulated. What remains is to fill out the details. The book raised hackles among the scientific community, perhaps because they were found out and did not want to be put out of a job, or perhaps because they knew how much work still had to be done in their fields and Horgan mistakes the limits of what is known to him for what is knowable in general. Central to Horgan's thesis is the claim that, unlike forms of creative expression like literature or art which are fueled by inexhaustible imagination, science inevitably runs up against the boundaries of physical reality. The laws of nature are not infinite in number and in the short course of a few hundred years science has found them all out.

Horgan, fatally in my opinion, approaches this without any reference to the theory of ignorance, which is the correct reference point by which to frame this problem. What does epistemology have to say about the end of science, specifically with respect to the theory of ignorance?

Supposing that the principle of uniformity of nature is true, the laws of nature hold everywhere in the universe. Regardless whether the universe is finite or infinite there is a fixed set of laws which describe its regularities over all spacetime. Ostensibly, this principle is not knowable with certainty, although plausible. To confirm that all the laws of nature have been found one must confirm that there are no regions of the universe fundamentally unlike those around the earth in which experiments are run, hence the problem are induction. Such a claim is inescapably an inductive leap lightyears in length, given that the

observational experience of human observers has ranged over a pitiful fraction of the sum totality of the universe. To add insult to injury, all of these observations are limited to the observable universe, which is only the subsection of the total universe from which light has had time to reach earth. So not only is the sum of human experience a fraction of the observable universe but the observable universe is a fraction of the total universe. Which, it the multiverse theory is to be accepted, is but a fraction of the multiverse.

Looking over all this immensity, one is first struck by how much darkness and obscurity prevails. Untold worlds, untold possibilities, scattered in unfathomable distances beyond human awareness. The unfathomable depths of the universe know no bounds. We may be able to assign numbers to it, but the depths of the universe utterly out-scope human comprehension. If you doubt this claim, picture yourself walking 5,878,625,370,000 miles. That's the distance traversed by a ray of light in one year. What might exist within this span? Now picture taking a stroll across 93 billion light-years. That is estimate of the width of but the observable universe. Science may or may not end in the sense that it has discovered all the governing principles of matter, but scientific analysis and discovery has quite a bit of runway left.

In fairness, according to information theory and statistical mechanics, much of this vast expanse is informationally uninteresting or thermally stale. It is in broad strokes a highly plain distribution of space dust and cosmic rays and abyssal voids. In the scheme of it all however, it is overwhelmingly likely that the universe in all its vastness hides wonders and oddities beyond the pale of what any one human knower ever thought possible. Consider for instance BPM 37093, the mind bogglingly massive diamond in space that is 5 million trillion trillion pounds in weight and 2,500 miles across.[11] . Is that anything you would expect to see or predict from average human experience? BPM 3709 promises to be only one of the staggering possibilities afforded by the laws of physics which has no analogue to human experience.

It may very well be the case that the principle of uniformity is true.

---

[11] https://www.universetoday.com/9295/astronomers-find-a-huge-diamond-in-space/

Gargantuan space diamonds, after all, are perfectly consistent with the laws of physics. It's entirely plausible that all physical unknowns conform to and are configurations of the known laws of physics. So Horgan's argument at least has the possibility of being true. That said *knowledge of a rule does not equate to knowledge of the application of a rule*. Nor does it ensure knowledge of how that rule may interact in all circumstances and will all other facts and rules included. knowledge forms a system and the whole is greater than the sum of its parts. Unobserved phenomena may exist which would cause scientists to utterly revise their assumptions about a given set of results that were tied to a given set of observations. The problem of induction looms, how to generalize from finite knowledge to all facts, both observed and unobserved. It might equally be the case that there are phenomena at work which we do not even recognize within the framework of existing science because there is no place for them. Such a discovery would prompt a revision of existing science. We only detect, in a manner of speaking, what our theories permit us to observe. What about dark matter and dark energy? These perplexing substances (?) are quite literally defined as *unknown* forms of matter and energy, both of which happen to comprise the *majority* of the matter and energy in the universe. It is entirely plausible that dark matter and energy are simply inexplicable from current knowledge but would make perfect sense with an expanded or modified physics.

Science also deals just as much in details as in theories—arguably more so. As biologist Thomas Eisner[12] said to Horgan, "I liked your book…But I do think you're wrong in a fundamental way. You equate science with the search for great universal truths."Eisner is on to something. It can be sliced two ways. It comes down to a basic metaphysical distinction between universals and particulars. While it might be possible that the universals, natural laws, have been identified by current science, it is absolutely impossible that it has identified all the particulars. If ever there was a known unknown, this is it. Let's also reiterate that the grounds for inferring that all universals have been discovered supposes that it is possible to survey the full scope of all possible knowledge and confirm with

[12] https://blogs.scientificamerican.com/cross-check/was-i-wrong-about-8220-the-end-of-science-8221

complete omniscient certainty that no law remains undiscovered. It is equally plausible that we have only discovered those laws which conform to everything within our bounded experience, and so we receive the illusion that everything is explained.

To elaborate further, what semi-universals science does identify are not immutable, eternal metaphysical categories or Platonic forms. They are contingent, revisable, flexible explanations that could be more accurately described as approaching the universal. Meanwhile, to the extent that we continue to be able to explore the cosmos, to the extent that is that our technology permits us to probe it further, because if otherwise restricted it could be explored for a practically eternity, science will never run out of work documenting the uncountable variety of phenomena permitted by the physically possible.

Whatever the case, Horgan runs up against fundamental problems of knowability. Our own finitude and fragmentary relation to the totality of existence damns us to some degree of ignorance . Hence the absolute necessity for understanding the context of the surrounding unknown which knowledge strives to overcome[13].

Beyond all possible variations of knowers and contexts of knowing within the physical universe there is one which transcends them all. By necessity it is the only one which transcends the physical universe itself, the only conceivable one which requires no context or preconditions and which has no finite constraints: God.

It is only in the mind of God that ignorance has no reality and is wiped clean. Such omniscience banishes all relativities; it is the ultimate context containing all lesser contexts, an entirely consistent and coherent field of seamless omniscience. God has no epistemic dependencies, or if it does, it they are between God's own conceptions. The mind of God is the only conceivable thing in which all vagueness, all uncertainties, all ignorance is banished in a blinding radiance of perfect knowledge. For Leibniz, the very fact that the world had a perceptible

[13] For the interested reader, I respond further to Horgan's thesis in the appendix.

rational order was a sign of God's mind at work. Otherwise, why should the universe have any intelligible order at all? Why should anything hang together? The reason, held Leibniz, was that the world was scaffolded by the logical relations of God's thought and thereby united in a pre-established harmony.

By picturing how any given relative knower might in theory deviate from the epistemic ideal of the absolute knower God, we can derive a kind of imaginary intuitive method for understanding how epistemic context and constraint emerges out of cognitive finitude. By calculating that delta, it is possible to see how the further away from that divine limit the epistemic agent becomes, the more ignorance clouds up and crowds into their mind, the less and less the rational order of the cosmos is perceived. Contrariwise, as any given knower gains in inferential power and approaches that unreachable asymptote, more of that order shifts into view.

This line of thought about the difference between omniscience and bounded knowledge leads to an intriguing conclusion. It is impossible to assume that knowledge is complete without presupposing omniscience, because only an absolute mind outside space and time would be able to grasp all epistemic possibilities. More accurately, there would be *no* perception of possibility, because possibility is a byproduct of ignorance. The omniscient would merely see all that is, not what could be. As will later be explored in greater detail, the number of things any given knower is ignorant is uncountable. In mathematics, an uncountable set is by definition infinite, because it lacks cardinality—a definite size or equivalently a largest number, and so it cannot be put into a one to one correspondence with the set of natural numbers. The quantity of ignorance is not, however, a purely mathematical relation. It may indeed be the case that there is a finite number of unknowns (however one might arbitrarily define a discrete, countable unknown fact), but this set of unknowns would still be uncountable. Not because of any mathematical property of sets, but because to know how many unknowns there are presumes partial knowledge of them. Put differently, the fact of how much is unknown is itself unknown, is itself included in the set of unknowns and perhaps even its subset of unknowables. If Horgan is wrong, it is ultimately for this reason: the finite cannot encompass the infinite.

Trying to understand knowledge without ignorance is like trying to understand numbers without the concept of zero. Zero is the crux of the mathematical universe; it balances everything. It is a number among numbers, but unique among all numbers. Without it, out goes the entire premise of positional notation, additive identity, set theory and arguably everything else about math as it is known today. Nothing may come of nothing, but something without nothing is unworkable. Zero, like ignorance, is substantial. Similarly, epistemology without ignorance is like math without zero, it loses its grounding, pivotal element that positions everything else in an intelligible reference frame.

Whenever learning is involved knowledge flows out from of ignorance. How is this possible? On an almost basic metaphysical level it is almost like assuming something from nothing. This is far from the case, because ignorance is a something. That's because in information theoretic terms, knowledge and ignorance are made out of the same stuff. knowledge, as will be explained, is an ordered substructure within conditions of maximum information entropy which contains all possible structures. In writer George Luis Borges' Library of Babel story, there is an infinite library containing all possible books and text strings. Thus every possible statement of fact is included within the endless seas of gibberish. The trouble is a search and identification problem. The inability to locate the valuable subsequences within the infinite string sequence, in addition to the inability to distinguish fact from merely plausible fictions, destroys the value of the library completely as a repository of knowledge despite it being the container of all textually representable knowledge. Signal cannot be extracted from noise. Ignorance first knowledge acquisition instead proposes that knowledge is gained by the recognition of ignorance, and then the pursuant mobilization of investigative behaviors to close the gap as it is explored. By this account the knowledge-value of the library would increase in proportion to how much smaller it gets as obvious junk is removed from the shelves. When ignorance prevails it its like a reader who is unable to parse the surrounding context of false text from true text.

The substance of ignorance, which is a probabilistic fuzz of possibility-

uncertainty, is successively shaved down of excessive irrelevant detail, informational degrees of freedom, until what remains is stable knowledge. There is no jump here between two incompatible media. Rather, the same informational substrate binds both knowledge and ignorance-as-uncertainty. The act of knowing is merely a tightening up and constriction of information states so that the least probable are excluded and the most probable are included into the developing knowledge representation. At no point does the background information conditions change between a state of knowledge or ignorance; its form is just rearranged.

The subset of what is known absorbs new elements of the knowable unknown. Such absorption is not always whole and perfect. To know or not know is not a simple, absolute binary condition. There are partial, fuzzy, transitional stages, in which static noise still clouds and intersperses with a spotty signal. Concepts can be partially understood. Since the universe is only ever partially understood, all knowledge is to some degree fragmentary either referentially, in terms of the under-description of its target subject matter, or relationally, in terms of how knowledge of one thing does not tell us how that knowledge relates to everything for which it is a relevant factor. For example. Someone who deciphers a tablet of writing from a lost ancient civilization likely cannot infer the entirety of that language from the text. Their knowledge is incomplete referentially with respect to its object. At the same time, their knowledge of this tablet's meaning alone does not tell them that there is a well known ancient language which is of the same family and which could be studied to gain context. Thus it is incomplete relationally. So knowledge can be incomplete both referentially in terms of its target object and relationally in terms of all possible related knowledge subjects.

We have seen how Bayesian probabilistic reasoning and other AI models help explain the referential incompleteness of knowledge. Meanwhile the concept of an ignorance graph, as a variation on the data structure called a knowledge graph, helps illuminate the relational incompleteness of knowledge. So even if we fully understand a concept we might not fully grasp the meaning of that concept in terms of how it stands in relation to other concepts that aren't

understood or even how it relates to other concepts that are understood. knowledge is not solely referential to its object, it is also relational to all other knowledge, potential or actual, and to ignorance. For these reasons ampliative inference is an inadequate model of knowledge generation. It is not always possible to go from what is known to new knowledge in a perfectly uninterrupted chain. knowledge is limited in its powers of self-extension.

The gist is that knowledge grows out of ignorance, remains incomplete, and then a seed of ignorance grows within that knowledge, prompting a further expansion and an additional iteration of the cycle. knowledge of a fact always depends on other facts connected to it within the nexus of causality. It's fair to say then that within knowledge typically resides a kernel of ignorance, symbolizing its internal incompleteness both referentially, in terms of what it specifically is about, and relationally, in terms of all external facts that condition what it is about which are not contained in the scope of that knowledge, and which would require learning an additional fact. Because even if something is known, that does not mean it is known what to do with that knowledge, or how it stacks up in comparison to all else. So ignorance resides in knowledge, until, of course, we find out about the limits of our knowledge by becoming aware of the flaw lurking within it. Once that awareness has dawned, then a kernel of knowledge begins to grow within the previous kernel of ignorance which expands as we become more aware of it, and so on it goes recursively, if not indefinitely at least until some ultimate terminus is reached and all relevant ignorance—or our awareness of it— has been zeroed out.

While there is a double incompleteness to knowledge, there's a unity to ignorance. What knowledge we do have is incomplete in the sense that it is not fully interpreted in all its internal implications. Knowing-that does not necessarily translate to knowing-how or vice versa. Nor does nor does knowledge of one fact supply the connective tissue that allows one to bridge that knowledge to another piece of knowledge or the discovery of an unrelated fact through ampliative reasoning. Ignorance, again, is monadic in its unity. Until knowledge recognizes it, the unknown unknown is a seamless, undifferentiated whole. Whatever knowledge we have exists in defiance and contrast with respect

to the unity of the unknown.

In the most basic terms this recursive dialectic between knowledge and ignorance captures the transactional exchange that is engaged between them during the process of learning. How else could knowledge appear where there was once ignorance, or ignorance erode and replace unfinished knowledge? This reciprocal, Russian-doll like concept, in which knowledge embedded into ignorance and vice versa, also encapsulates the key insight that knowledge and ignorance are not diametric opposites. They are deeply intertwined and interpolated. This is a valuable insight because until the conventional image of knowledge and ignorance as polarized, one to one symmetric, reversed mirror opposites or binary logical complements is overcome, epistemology will never be on the right footing. The implicit informational unity between knowledge and ignorance is obscured by this inaccurate and simplistic is/isn't distinction.

In what has since become an integral part of ignorance studies lore, in a 2002 press conference the then U.S. Secretary of Defense Donald Rumsfeld made an astute offhand remark about the possible relations between knowledge and ignorance. In regards to the question of whether Saddam Hussein was in the possession of weapons of mass destruction, Rumsfeld made a distinction between known knowns, things we know we know, known unknowns, thinks we know we don't know, and unknown unknowns, things we don't know we don't know.

This unexpectedly philosophically brilliant statement pithily summarizes the set of possible conditions involving relative degrees of knowledge versus ignorance. It depicts a compact factorization (known-knowns, known-unknowns, unknown-unknowns) of each combinatorial pairing. (After Zizek, DeNicola (2018) proposes adding a fourth category, unknown knowns, to the list. This inclusion to my mind is controversial, so it will be dealt with separately in the next following section.)

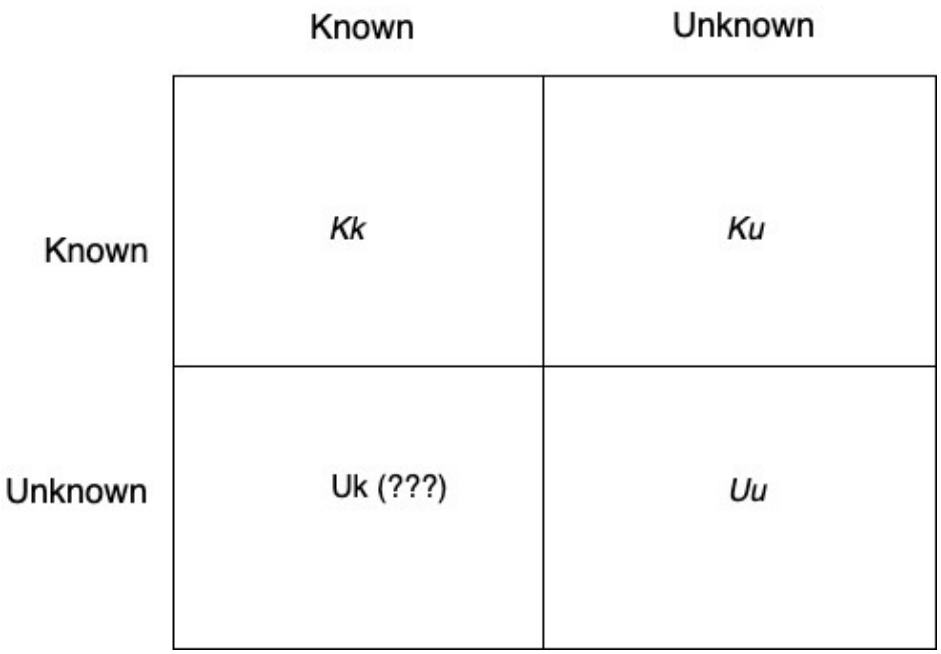

The first category (known-knowns) is meta-knowledge. Meta-knowledge has a critical indexing and accounting function. By knowing what is known the present state of knowledge is self-registered and confirmed. In an additional sense this self-accounting also indirectly nets you knowledge of what you don't know, at least in broad strokes, because once you know what you know you'll be able to detect when some new finding conflicts with it, thereby signifying an unknown. However, the matter is complicated by the fact that the boundary between known-unknowns derived by evaluating known-knowns is fuzzy and poorly distinguished from the unknown-unknown from which the known unknown juts out. The subset of known-unknowns ultimately shade off into a superset of unknown-unknowns, where total ignorance prevails and no particulate facts can be discriminated from the featureless, blank of ignorance in its entirety. The fundamental question then becomes: How to convert unknown-unknowns into known-unknowns and then finally into known-knowns?. This happens inevitably through either accidental or intentional learning, but it's wrong to assume that how it happens is obvious.

There are a few approaches to penetrate the unknown-unknown. The first means is not a method at all. It is simply accident, a random walk through the space of possibility that leads to bumping into surprises and unplanned discovery. It is having things happen to you and learning about them after the

fact. By putting oneself out there with a vague understanding that you will eventually be exposed to new information, the unknown-unknown can be gradually sounded out and brought under cognitive control. Or, less gracefully, you can simply stumble into a fact. This is by no means effective. What it gains in generality it lacks in specificity: you are guaranteed to learn something if you just expose yourself to new experiences, but it may not be useful or what you wanted. Establishing a bridgehead into the totally unknown can also be achieved by following where known-unknowns lead. This of course begs the question: how does the known come to reference the unknown? The short answer, I would suggest, is metacognition. By registering and appraising what you know you know, any new stimulus that declines to fit one of those categories must then be a known unknown. Thus the unknown becomes a recognizable object relative to existing knowledge. It is possible to infer to some extent, what one doesn't know from what one does know. Such known unknowns ultimately taper off into and terminate at the unknown-unknown. Thus, through a process of investigating known-unknowns, the broader unknown is gradually revealed. Through successive waves of investigation the unknown is drawn into the known. Separate pieces of knowledge, at first seemingly unrelated, isolated islands within a sea of ignorance, gradually emerge out of that ignorance and unveil an intertwined mosaic. Tugging on a single piece, like pulling on an exposed link in a chain buried by sand, uncovers that they are in fact all connected and move together as one.

knowledge progresses by flowing back and forth between these stages, from known known, to known unknown, to unknown-unknown, back into the known unknown and then the known known. During each iteration step, the territory of the unknown-unknown comes in view of the known through the intermediate linking step of the known-unknown. The mechanism that describes this process we call inquiry, though here is not the place to take this matter up.

This terminology is an indispensable shorthand for referring to knowledge ignorance relationships. Without it, it's difficult to distinguish the gradations that separate knowledge from ignorance.

The concept can also be extended to address modalities of knowledge, such as

whether it is possible to know something. Understanding these modal relationships are much more complicated however. Rather than to delve into them here, and given that part four of this book is dedicated specifically to modal reasoning, I've decided to include it in part four than pile on too many similar concepts in one place.

There is more one combination of these terms that has been left out, the unknown known. First suggested by social philosopher Slavoj Zizek, it is something of the black sheep of the four, defying the smooth transitive logic flow proposed above. These unknown knowns are the submerged unarticulated assumptions that govern our activities, the tacit knowledge which is effectual without being recognized. The riddle of unknown knowns is just a window into a deeper enigma of the unconscious mind and deserves its own separate treatment.

It's customary to assume that conscious thought is primary. To think in the fullest sense of the word is to be think consciously. Indeed, an unconscious mind is seemingly a contradiction in terms. Consciousness is the paradigm of mindedness. According to the axioms of epistemic logic, to know means one knows one knows. Of all properties consciousness is unambiguously what sets minds apart from non-mental things. Unfortunately for common sense, science has done little to support the intuition that thought is mostly if not all conscious. Quite the opposite, converging evidence suggests that there really is an unconscious mind, and if anything it is far more industrious than conscious thought, which is more like window dressing. (Philosopher Peter Carruthers even goes so far to argue conscious thought does not exist! [https://pdfs.semanticscholar.org/7632/4513e02e352b08349f2e4e511a578f67e89c.pdf]) Most thought, or at least information processing—assuming there is a difference—happens unconsciously ([https://www.researchgate.net/publication/223488264_The_cognitive_unconscious_An_evolutionary_perspective]). Studies have even shown that decisions are made up to 10 seconds before the subject is conscious of having decided ([https://www.nature.com/articles/nn.2112]) Even with all this data, the existence of an unconscious mind poses a

vast philosophical quandary. It would mean that the very engine of knowledge, the mind, is shrouded in self-ignorance. The seat of knowledge rests on a bed of ignorance, its own mechanisms of knowledge production inscrutable to itself.

Such a prospect of the mind being unknown to itself, if it did not originate with Freud was cemented into popular culture by him. The mind, thought Freud, receded into cavernous depths of unconsciousness. Stashed within this darkened subterranean cave complex were those truths about ourselves we refuse to acknowledge, repressed traumas, ugly desires, simmering resentments and all the ghosts and ghouls that torment the human psyche. Like invisible poltergeists knocking down glasses in the cupboard or flickering the lights, these unconscious forces bedevil us without making themselves explicit, diminishing our quality of life and harming our relationships without us being aware of their nature and influence. Psychoanalysis, Freud held, was the method of surfacing these pathological maladaptive unconscious mental patterns to the conscious mind and vaporizing them in the white heat of rational self-reflection. The idea that the mind was not perfectly transparent and translucent, but compartmentalized into hidden regions, dead ends and blocked off egresses, placed a black hole of ignorance where the sun of reason once stood. Irrationality, psychic anarchy, instinctual drives, were the law of human nature. Descartes' *cogito*, the perfect certitude and clarity of conscious rational self-awareness, was instead rife with illusions.

A natural criticism of this idea of unconsciousness and also the more modern ideas that view it as an engine of information processing rather than a lockbox of repressed emotions, is to wonder how to validly attribute properties to an object which by definition is hidden from any kind of direct inspection. At what point is a supposed unconscious process more like a mechanistic reflex than a mental event? What is the design rationale for this dual partition of the mind into conscious and unconscious departments? Why does consciousness exist at all, if the unconscious mind has such abundant functionality? Why are we not "philosophical zombies" if consciousness is merely an afterthought or epiphenomenon? If the unconscious mind is involved in blind information processing, does this give us the right to ascribe mindedness to computers? Are

they not information processors? How does consciousness emerge from it? It is difficult to exhaust all the questions.

Because there is no means to inspect the unconscious mind, anything can be ascribed to it without the possibility of clear confirmation or refutation, and this is even true of the *very idea* of an unconscious mind. In other words what does not actually exist and what is unconscious stand on the same plane and are indistinguishable; something that cannot be shown to exist is evidently no different than something merely presumed to exist or does not exist. The conscious mind, in contrast, is utterly apparent—it *is* that which is apparent, at least introspectively. Though its nature may be mysterious, consciousness has a self-asserting reality to whomever experiences it. It is Descartes's cogito. If I am aware of anything I am also aware that I am aware. The unconscious mind, much like the unknown writ large, has no clear boundaries, it retreats into bottomlessness.

We can do things without knowing how we do them and experience without knowing how experience works. Such loopholes and blindspots seemingly contradict the view that there's a transparent and uncomplicated relationship between the knower and the known. When the object of knowledge is themself, knowledge suddenly becomes ungrounded. In a manner of speaking there's no "room" for such understanding; the modules and structures of the mind do not self-referentially process their own processing. If the mind is a representational engine, there is no reason it should come with the ability to represent itself to itself right out of the box, not any more than a camera should take a picture of its own internal mechanisms when it snaps a photograph, or as in an M.C Escher drawing, a hand with a pencil should draw itself drawing itself. Such self-representation is similar to a von Neumann universal constructor, a self-reproducing machine the output of which is an exact replica of what produced it. In this case the copy mechanism is itself copied, just as in the case that self-knowledge represents what represents. There is an essential difference however between these comparisons. In a von Neumann universal constructor, the description tape which defines the replication mechanism is held separate from the mechanism. The faculty of knowledge, implemented as a global property of

neural circuitry (in all likelihood there isn't a single local region of the brain that deals with all knowledge) is designed to describe things that are not itself, and "itself" has no stable persistent existence anyway. Indeed it should come as no surprise that much of how the brain works and problems related to conscious or unconscious processes are so hard to understand. Like the eye that cannot see itself, the center of knowledge, where all information converges, is not constructed out of what it constructs. The eye is not made up of vision, the knower is not made up out of knowledge.

The map/territory distinction breaks down completely at this point. There is an expectation that the map is external and derivative of the territory which allows the map to describe it without constituting what it describes. This discussion touches on the problem of cognitive closure that will be expounded in detail elsewhere. Are there unthinkable thoughts? Unknowable truths? Specifically with regards to the mind's understanding of itself?

Does such an inner contradiction within the knower in relation to itself destabilize or undermine the integrity of what knowledge the knower does have? The most obvious complaint is that such a flawed knower goes against the vision of one empowered by faultless reason, reason which can penetrate all facts equally. Knowledge happens within a plenum of knowledge, a seamless context of knowledge without gaps or contradictions between its logical links. Among all the features assumed of a knower, one of the most essential is the premise that its basis for knowing must be *consistent.* The unconscious mind puts paid to the notion of the positive introspection axiom, often formulated as:

$Ki\phi \Rightarrow Ki\ Ki\phi$

$(Ki(A) \wedge Ki(A \rightarrow B) \rightarrow Ki(B))$

(http://www.sci.brooklyn.cuny.edu/cis/parikh/omnisci.pdf)

Which reads: if i knows ϕ, then i knows that it knows ϕ. Such self-consistency implies an unbroken chain of knowledge and knowledge of knowledge which assumes a property of *closure*: logically airtight epistemic conditions in which any inconsistency— another word for ignorance in this case—is banished. knowledge frictionlessly begets knowledge of itself without disruption. In this case however, i may unconsciously know (or think) ϕ, and therefore lack the

salience to derive that it knows it knows it. This is a problem because the axiom of positive introspection is supposed to be a fundamental rule of epistemic logic on which all valid propositions are founded. Tacit unconscious knowledge, being inaccessible to introspection, interrupts the logical chain. Ignorance it would seem looms larger the closer it is to the knower's knowledge. One would hope that the knower serves as a special boundary condition that excludes it from the envelope of ignorance, after all, it is what cuts through ignorance like a laser, and sublates it into itself as knowledge. Ignorance ought to fade the closer it approaches the knower rather than enlarge, as it is assimilated into the consistent universe of knowledge. What underserves knowledge should itself be known, otherwise, knowing rests on ignorance and is exposed to an inconsistency. We do not live in such an ideal epistemological universe.

When one steps outside the sometimes restrictive confines of logic and looks at how the brain is constructed and encodes knowledge, these muddles can be set aside. As soon as the demand for perfect rationality are cast aside the underlying reasons for irrationality are ironically brought to light and made reasonable. There's a method to the madness.

The Society of Mind hypothesis proposed by Minksy holds that the mind is composed of encapsulated competencies that each excell in a specialized task (like visual edge detection or parsing the phonetic components of speech) but none of which contain the capability of the whole. If this account is true it would make sense that nowhere in the mind does it "all come together" as a cohesive self-aware unity. No such Cartesian theater, as Dennett calls it, can justifiably be said to exist once you consider the actual composition of the brain as a modular architecture ([https://ase.tufts.edu/cogstud/dennett/papers/tempanom.htm]). The conscious mind can be likened to a city newspaper that captures the most important headlines of what's going on rather than the routine activities like garbage collection or parking citations, signaling what is exceptional rather than humdrum (Eagleman, 2011). Brain function might be likened to the software concept of abstraction, in which interdependent modules interact only at their entry-points and end-points but do not rely on assumptions about how the other modules work or what they do. The brain solves many small problems rather

than one big problem. The riddle of its functioning and the aforementioned inconsistencies of unconscious knowledge appear intractable because there is indeed nothing implemented globally and only many interacting subsystems operating in parallel.

Any knower that can be said to inhere within this network of activity is therefore ephemeral, always shifting and flowing as various systems come online or go offline. If we don't know how we know it's because only some of these subsystems are active at any one time and the knower as a complete self and subject of knowledge is an illusion. These various subsystems are occasionally recruited into processes and at other times go offline, never retaining a uniform cohesiveness of identity. Contra Descartes, there is no integration center where it all comes together as the nucleus of the thinking subject. The pineal gland does not do what he supposed it does. Perhaps the best candidate for such an integration site is the thalamus, which is like a switching station that relays and coordinates incoming sensory signals and cortical activity ([https://www.cell.com/neuron/pdfExtended/S0896-6273(20)30005-2]). Even then, the thalamus's physiological role is mediated by its interaction with transiently activated circuits distributed throughout the cortex ([https://www.frontiersin.org/articles/10.3389/fnsys.2019.00043/full]). Its purpose is to connect and associate operations throughout other regions of the CNS, it is a gateway through which a slipstream of signals flows. The more basic properties of consciousness depend on arousal and awareness which are associated with more primordial structures in the brainstem ([https://n.neurology.org/content/87/23/2427.short]). Consciousness does not happen "in" the thalamus or the brainstem.

The confusion about unknown knowns therefore originates from the false requirement that knowledge in real life must conform to the idealized logic of knowledge worked out in purely formal symbolic terms. knowledge is processed at a sub-symbolic, subliminal level of unconscious processing, which does all the heavy lifting. No single part of the brain is responsible for knowledge, and one can reasonably infer that all the parts that do are themselves not conscious, and themselves not knowers. Though it is extremely difficult to figure out how these

neural networks integrate their activity and resonate together to yield knowledge, that does not invalidate the reality of that knowledge. It's not necessary to know how one knows to know. Unknown knowns underscore the deep interrelationships between knowledge and ignorance and shows that even the means by which knowledge is achieved rests on a foundation of ignorance. This is not a problem, however, because no matter the situation the end result is still knowledge.

As the examples of AI research in the previous chapter demonstrate, epistemology can learn much from technological innovations in machine learning and related fields. Too often this research remains boxed up in its technical and applied domains, but much is to be gained from breaking these findings out of the original frames in which they were cast and subjecting them to broader interpretation.

Among these one most pertinent for current purposes is what are called knowledge graphs (KGs). KGs are graph-theoretic data models that store entities in terms of nodes and edges representing contents, and semantic relations between contents respectively. The idea kicked around since the middle of the 20th century, but it only took off in stride with the rise of the internet and specifically with the rise of the Semantic Web in the 2010s which strove to make data on the internet machine-readable.

To get a basic idea of how knowledge graphs works, consider a simple graph captures facts about Mozart and the "region" of entities closely associated with him. Mozart might himself sit in the middle of this network, the spoke in the wheel. Extending from this Mozart node would be various edges standing for relations, such as "born in" which in this case would connect to "Salzburg", or "known for" which would connect to "classical music composer." The node "Salzburg" would in turn have a relation "city of" or more abstractly "belongs to" that connects to "Austria", sharing this connection with Mozart who has the same endpoint by way of the relation "has nationality of." Austria would then connect to Europe, and so on it goes, extending out into the countless variety of interconnections and entities that make up the relational semantic network

holding between objects, concepts, and facts in the world.

Advances in artificial natural language processing have resulted in the construction of vast knowledge graphs from oceans of online data. Notably in 2012 Google debuted their own graph which models the associations between millions of entities and which more or less has come to be considered *the* knowledge graph. knowledge graphs have many useful applications, from enhancing search results, prediction engines, fraud detection, and more. These applications however are not the present concern. The actual implementations of knowledge graphs in computing are only of secondary interest. The main goal is to see how these concepts can inform the understanding of knowledge-ignorance relationships widely construed.

One of the most striking properties of conventional KGs is their overtly positivistic structure. They are closed off to any notion of ignorance and it has no standing in the semantics. They encode networks of positive facts. Their design tacitly takes it for granted that the world decomposes into something like unambiguous, static atomic facts. Each subgraph is a tightly wound molecular bundle of truth about an object and its relationships. There can be no degrees vagueness to the manifest relations asserted between nodes, and the nodes themselves are not analyzable into more elementary primitives. Especially given that in some simpler knowledge graph architectures, a node is simply a proper noun, a literal string of text without any deep content. And its relations are data mined from syntactic regularities parsed from countless sentences scraped from sources like wikipedia. The data dictates what the graph knows, but it's safe to say it doesn't know that it knows it. It blindly deduces the structure of relations inherent to the data flows that have been etched into the contours of redundant internet traffic and distributed among its boundless reams of hypertext.

Plainly, it makes no sense conceptually to conceive of pure "ignorance graphs" as any set of unknowns lacks determinate relational structure. Indeed, the truly unknown would not even be quantifiable as a countable set, it is totally undefined. The nodes would be undefined with no relations to assign to them. Or rather, there would be no discriminable nodes at all. In technical terms, considered as a KG, an ignorance graph would be considered a null or order-zero

graph. The situation changes when we consider mixed knowledge and ignorance graphs. To introduce a concept of ignorance we have to posit a concept of learning and growth to the KG, conceptualizing it as propagating over a data set which, prior to processing, constitutes a sort of field of ignorance. The graph probes its "environment" for tentative, conjectural relationships, grading each discovered node based on their consistency with the preexisting graph. In other words an initially unstructured, growing graph associates nodes in the course of learning by tying them in to it, and based on how these relations stack up, is able to "discover" nodes which were originally not part of its network.

A data model can indeed represent unknowns so long as they are "tagged" by a known relation, i.e, they are known unknowns. A graph that starts off small and grows, must traverse a space of unknown nodes, which are as it were latent in the unprocessed dataset. Ignorance nodes therefore are still embedded in the space in the data in the graph develops, only it has yet to traverse these nodes and assimilate them into its graph.

As it happens there has been research in the knowledge graph space that address these problems. One 2017 paper in particular *Marrying Uncertainty and Time in Knowledge Graphs* confronts these issues head on. (Chekol, M., Pirrò, G., Schoenfisch, J., & Stuckenschmidt, H. (2017). Marrying Uncertainty and Time in Knowledge Graphs. Proceedings of the AAAI Conference on Artificial Intelligence, 31(1). Retrieved from https://ojs.aaai.org/index.php/AAAI/article/view/10495)

Traditional KGs, the researchers point out, lack temporal and probabilistic dynamics. This is a glaring flaw considering that "the majority of facts are fluents," meaning that they are time-dependent and changeable. Such systems lack "inference or deduction rules and consistency checking constraints \[that are] useful to both derive implicit/new facts from existing ones and constrain or identify conflicting facts", (p. 88). The researchers proceed to develop temporal and probabilistic graphical models to account for these deficits.

What is missing from their analysis however is a coherent concept of an "ignorance node" as a logical atom. It is clear that the nodes of the graph that are uncertain, nevertheless contain a hidden true value, and that the system must

engage in noise reduction to arrive at this value. In a state of uncertainty, the indeterminate value is still a node, even if it has not yet been evaluated. The uncertainty is a concrete fact about the data, not merely a reflection of the untuned model.

If positive knowledge can be thought of as a web of true propositions related by edges, each depending on each other for consistency, ignorance-as-possible-knowledge is those nodes that are existent true propositions but unconnected to the graph but filling out its data environment. These nodes are not nothingnesses; but potential knowledge that has yet to be assimilated.

One can abstract a KG so that, independent of the values of the terms, there is an associated property that defines which knowledge-ignorance predicate the node represents. For example:

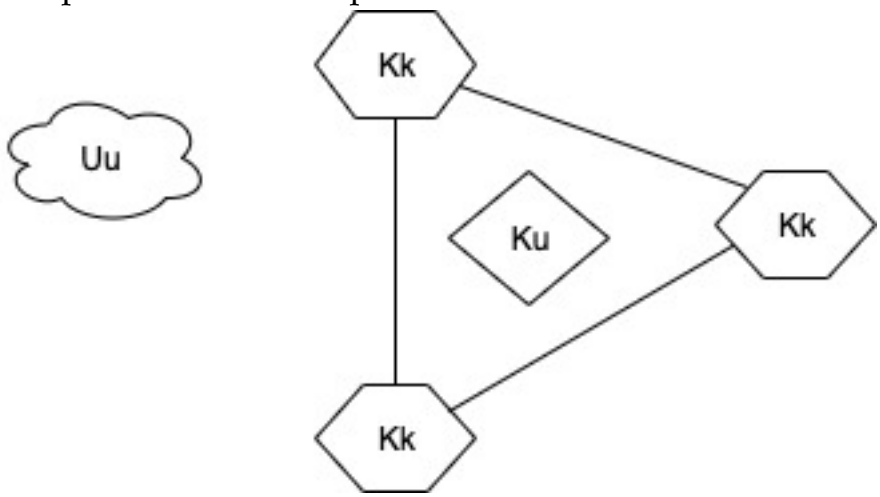

This diagram demonstrates a toy model of a graph that illustrates relationships between known-knowns (Kk) known-unknowns (Ku) and unknown-unknowns (Uu). Through some otherwise unspecified introspection operation the Ku contained within the bounds of the Kks is added to the graph:

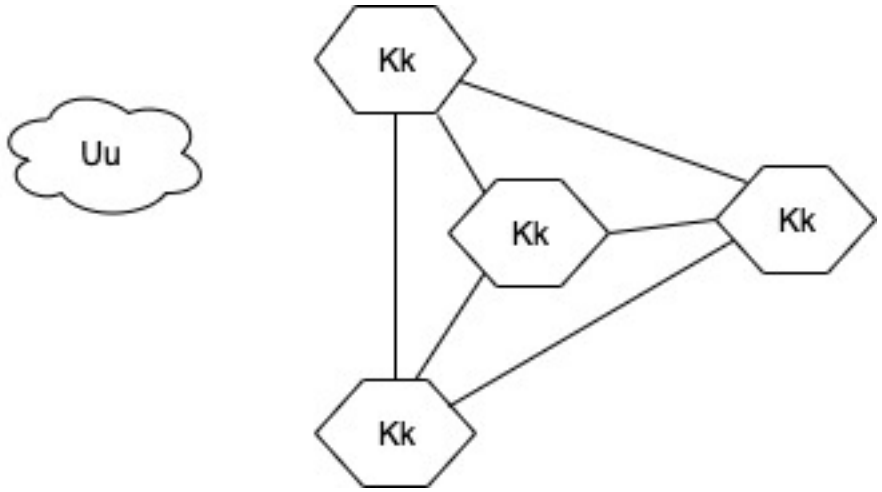

This operation happens at a given time interval during which the graph iterates. So each state of these diagrams represents an incrementing time interval of t + 1. They are meant to depict how known unknowns are contextualized by the context of knowledge.

Meanwhile, because there is not pathway to the Uu, it remains outside the reach of the graph. There is no opportunity to associate with preexisting knowledge it so it is unrecognized. Another way to put it is that all the necessary preconditions for updating the state of Ku to Kk are satisfied. Uncertainty about it is eliminated as it is probed and subjected to observation. However for Uu, no such preconditions have been met.

For this model to be truly interesting there would have to be a definition for the relation operation that associates Kk nodes to a Ku node. One might consider some inspection mechanism along the lines of a Bayesian check which computes whether the Ku matches the priors of a weighted graph. Moreover, the graph exhibits both an inspection mechanism for what it encloses and an observation mechanism for nodes that fall outside its enclosure.

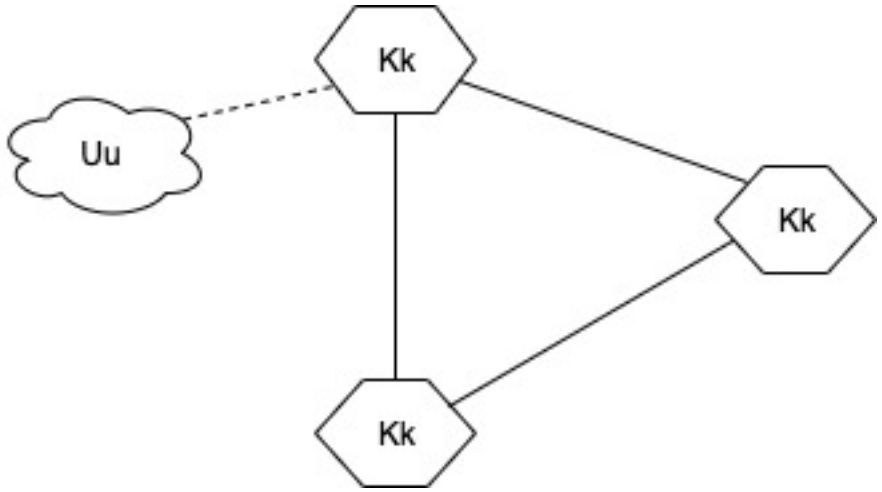

One might suppose that each node extends a tentative "observation relation" to local external nodes outside the graph and evaluates it according to some inference mechanism. The observation relation is analogous to an act exploration or even curiosity. This mapping of the observation function, for our purposes, can be randomly assigned at first to any "local" unknowns near a given Kk node. This is to represent undirected environmental exploration and the fact that certain knowledge brings one closer to certain possible discoveries conceptually "adjacent" to it. Once an external ignorance node is evaluated through observational inspection, other regional Kk nodes extend a relation to the same node. The graph "sees connections" between different parts of its knowledge base and the external node.

Furthermore suppose that we can extend the rules by which a Ku can be upgraded to a Kk based not just whether it is enclosed by Kk nodes but also whether an observation function successfully evaluates it to true. For simplicity let's label the observation function f(Uu):

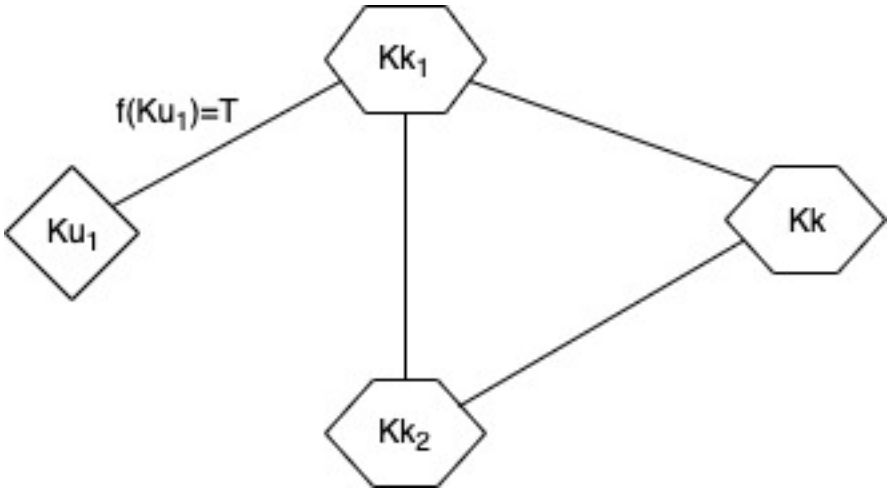

If "enough" (again, subjective) nodes connect to the same external node, the node is upgraded to a Ku.

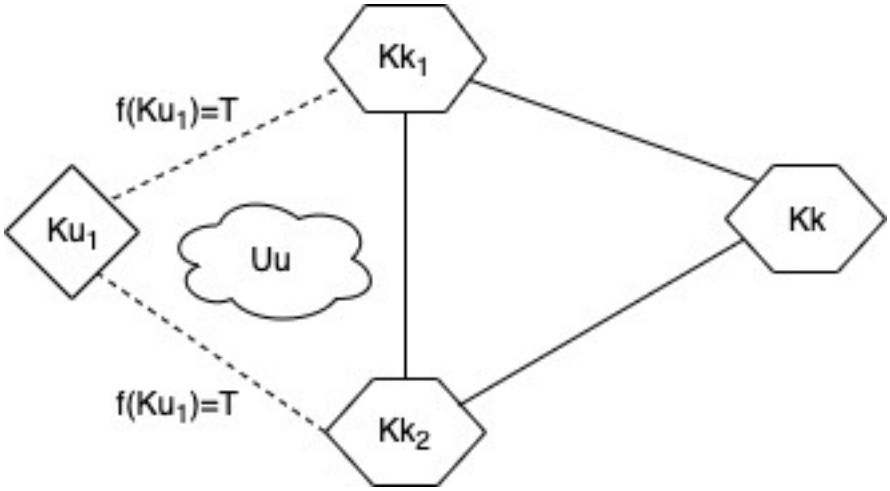

Let's assume that another node Kk2 then "picks up on" the fact that Kk1 has successfully evaluated Ku1. Notice also the addition of another unevaluated Uu which simply did not get discovered before the previous one and so falls within the potential scope of Kk1, Kk2, and Ku1. Kk2 repeats the same series of operations as Kk1 and establishes a confirmed connection to Ku1 as well:

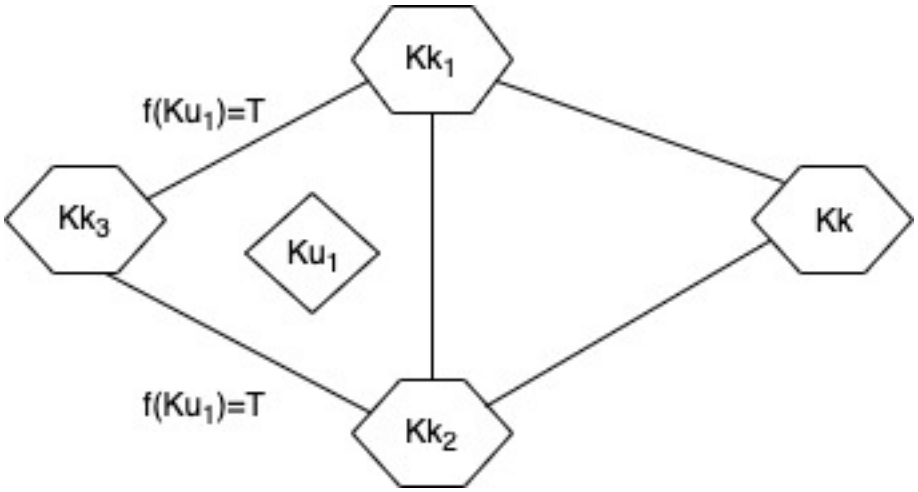

Applying the rule of enclosure, the additional Uu is upgraded to a Ku.

Nodes falling within the enclosure, are preferentially evaluated before nodes falling outside the enclosure. One comes to know what is contextualized by existing knowledge before learning something completely unfamiliar. The items that are introspected are added to the knowledge graph before items that are observed:

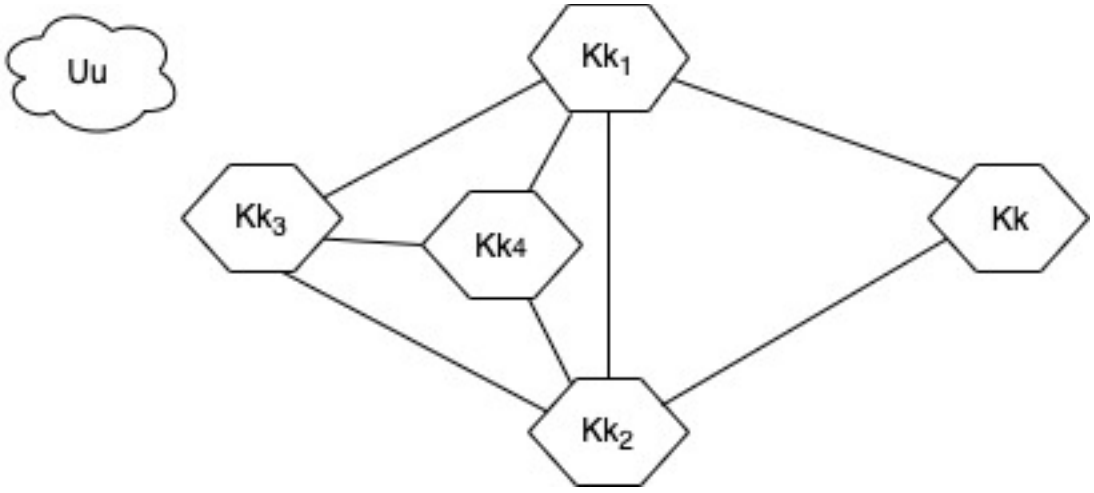

One must also consider that unknowns-unknowns are not sparsely distributed but replete and pervade the graph's outer environment:

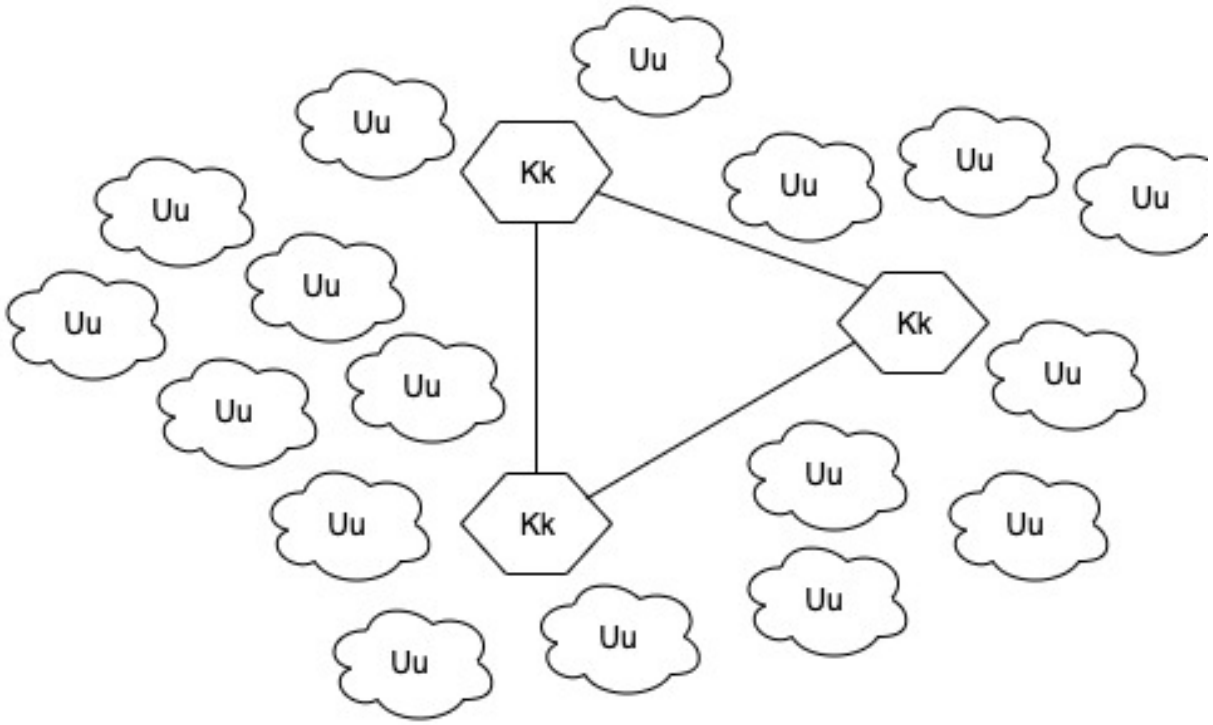

While these rules are admittedly unformalized, they don't have to be to demonstrate the key points. Namely, that knowledge graphs can be broken out of their technical framing and applied intuitively to understanding knowledge-ignorance relationships. Additionally by incorporating time and iteration steps into the picture, we are able to derive a concept of ignorance nodes as the as of yet unprocessed data points within the data set.

Caveats about the lack of formal rigor aside, I like this visualization because it it depicts knowledge as a relational network that iterates and expands over a set of unknowns through context. It illustrates how unknowns go through a series of transformation steps, from Uu to Ku to Kk in the course of its sublation into the knowledge base. To learn anything new it must either be discovered by accident and evaluated or induced form previous knowledge, or fall within the coordinates of what has been learned and deduced.

In the book *Superforecasting: The Art and Science of Prediction* political scientists Philip Tetlock and science writer Dan Gardner examine the principles of forecasting: how to anticipate future events. They report that, according to Tetlock's research, self-professed expert pundits are only marginally better than

chance and hardly better than those of non-experts at predicting major events. Some, however, are better than others. Tetlock found that these skilled forecasters had several traits in common. They did not trust in loose intuitions or gut feelings but invariably relied on an arsenal of quantitative methods. Such *superforecasters* emphasized preciseness in their estimations and had a high tolerance for uncertainty. They ruthlessly subjected their models to criticism and strive to prove themselves wrong. They are "dragonfly-eyed", capable of integrating multiple perspectives and correlating patterns across several divergent threads and associating seemingly disconnected events related by hidden interdependencies rather than narrowly focusing on the single obvious domain of import. Borrowing a phrase from software development, Tetlock describes them as in "perpetual beta", a state of continual improvement and adjustment that grows with the uptake of new information and benefits from error by correcting the model to accomodate for that exception it in the future. Their models are flexible rather that rigid: they do not expect reality to conform to their expectations but rather the opposite. Above all, they are not averse to uncertainty. They are at home with it and remain humble before the complexity of the world. They refuse to lull themselves into a comfortable, false sense of certainty but are perpetually on guard against the prospect of being wrong and thus actively search for reasons why their models might be imperfect. In essence, superforecasters are Bayesian to the core.

Here it helps to take a step back and look at what a forecaster is trying to do with what knowledge they have. Tetlock does not conceive of the challenge of prediction as a rivalry between knowledge and ignorance, even though it is quite implicit in the concept of forecasting. Forecasting is the mobilization of incomplete knowledge against unbounded ignorance. Predictive knowledge is hampered by the constraining effects future-oriented information incompleteness. Such forecasters, like all of us, are unable to glean information about the future, and run up against the problem of induction using present and past information as a tool for scrying it. Instead, they play a game of chess against themselves, creating a model that duels with itself, attempting to discover its own shortfalls before the future comes to pass and does it for them.

To the limited extent that ignorance can be conceptualized as potential knowledge, knowledge and ignorance exist on the same grid. In this respect the tension between knowledge and ignorance can be likened in a revealing and nontrivial way to the old board game *Battleship*. In *Battleship*, two players set up game pieces on their own two dimensional grid. Each point on the grid represents the intersection of a column and row, such as A-1, B-2, etc, the letter specifying the column and the number the distinct position along a row. Each of the game pieces, the battleships, is of a certain extension in length, and can be positioned on the grid either vertically or horizontally and thus it occupies a certain set of coordinates along it. Neither player can see the other player's placement of pieces on their grid. Rather, they take turns "launching torpedos" by specifying a point on the grid where they think their opponent's pieces might be. As the game progresses, the players track where their hits and misses land on a separate "tracking grid" that allows them to record the history of their torpedo launches and thus map out where the opponent's pieces are by narrowing down the available spaces on their grid.

In a way the humble game of *Battleship* is the quintessential analogy for the antagonism between knowledge and ignorance. (See Van Ditmarsch 2002 for a general overview of such so called knowledge games, which are arguably just as equally ignorance games [https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8586.00134]) The entire premise of the game is to map out the ratio between known and unknown by successively approximating the true layout of the board through a series of hits and misses, in other words, by mutualizing internal and external information. The misses, while not as informative as the hits, nevertheless convey information and enhance the future probability of a hit. Unlike real world scenarios, because the grid is finite in its dimensions and the placement of pieces are static, the proportion known and unknown can be exactly quantified in the same way the fullness versus the emptiness of a glass can be. The unknown, starting at a state of maximum entropy, is whittled down by the known. In this way *Battleship* encapsulates the antagonistic relationship between knowledge and ignorance. What knowledge we do have is purposed to target what ignorance it selects as pressing, strategically important, significant

and task-relevant. We use knowledge as a weapon against whatever outlay of ignorance our awareness of discloses.

In healthy epistemic conditions in which one is not delusional or misguided by positive ignorance, an epistemic agent's perspective is almost one of combative hostility toward what they do not know but seek to know, or vigilance towards "surprise attacks" on their plans launched by the unexpected. The healthiest attitude one can have on ignorance is that it is a kind of target or risk. The knowledge seeker can be said to be in a perpetual struggle with ignorance. Warlike metaphors are not so untoward given the dangerous obstacle that ignorance so often is. And the more aggressively we pursue the conquest of ignorance, the more we are aware of its scope and danger, the more it motivates us to defeat it, the better we are positioned to achieve our goals. Indeed the more challenging the learning goal one sets, the better the performance (Latham, G. P., Seijts, G., & Crim, D, 2008). It is by challenge and struggle with ignorance that knowledge is gained and through which practical knowledge is tested and exerted. The more that we are confronted with the enormity of our ignorance, the more we are compelled to reduce it.

People naturally conceive of ignorance as an antagonistic force. Plato said the root of all evil is ignorance. If anything is to spoil our aims, to frustrate and defeat us, it is ignorance. Else we'd simply reach out and take what we wanted by following the course of action that is plainly certain to work as assured by our knowledge. If knowledge is power it is only because ignorance limits us most. Learning is therefore a kind of epistemological self-defense or adaptation, a method for building intellectual fortifications and for launching attacks on ignorance and fending off its assaults. It is a method not simply to remove our ignorance as we would a pile of debris impeding our path, but to arm and train ourselves to be wary of making mistakes, to sense landmines of ignorance in the field ahead. As soon as we become conscious of our own ignorance and recognize how fully it inhibits our possibilities for effective action, we go on the offensive. It cannot be tolerated. Phenomenologically, knowledge of our own ignorance becomes an immediate irritation like a toothache that does not subside until the cause of it is removed, although the pain of such ignorance is

proportional to how much the missing information is goal-relevant.

Much as maintaining good epistemic health involves a kind of self-competitive match off between ourselves and our consciously apprehended ignorance, there is another form of ignorance that is neutral in character. Objective uncertainty, manifested perhaps, as we shall see, in the probability waves of quantum mechanics and other fundamental physical phenomena such as entropy, is also just a fact of life and a property of the world like gravity or heat. Gravity and heat can ruin your day (try falling off a cliff or take walk through fire) but these properties in themselves are not hostile. The value of it is contextual. At times we are comfortable not knowing, but that is because those regions of ignorance to not stand between us and our goals. Uncertainty can at times be playful, exciting, and useful. A pleasant surprise is more rewarding than an expected source of pleasure, and the thrill of gambling comes from the prospect of losing versus winning big. And one can be sure that it is to your benefit that a hacker is uncertain of your account password. It would seem that more often than not if we are uncertain whether we want to happen will happen, uncertainty is by and large stressful and better for it to be extinguished.

All that said, there are subtle shades at play which prove to be exceptions to the rule that one should always be on guard against ignorance. While it is wise to approach any new situation with the understanding that you probably don't know at least something about to it, the more expansive definition of ignorance we are working towards will show that there are even cases where ignorance can be constructive, strategic, and beneficial. On a deeper level ignorance can morph from a passive state to an active choice, as "to ignore" means to select what to pay attention to, or to "repress" means to deliberately shut down an unproductive line of thought or hide from a truth that has no practical benefit or could even be detrimental in knowing. In general however, the agent's relation to ignorance is adversarial. It is more often than not a roadblock to be worked around or to overcome.

The idea that learning is a kind of epistemic self-defense leads naturally to a

concept of uncertainty-proofing. Because uncertainty is in a sense permanently and pervasively *out there*—mapping to the countless non-repetitious things happening in the universe we don't know— the best we can hope to do is "cultivate our gardens" as Voltaire said, attend to what we can control and fortify with local knowledge. We must cling to our islands of fragile certainty and build out from them, but we must also fortify and defend them. We need to create regulated environments that are causally structured to be as predictable as possible and minimize noise.

Habit is a direct consequence of this need for uncertainty-proofing. William James, in his monumental *Principles of Psychology* dedicated significant attention to the study of habit. Habit, according to James, was more than a state that someone settled into out of cognitive inertia or custom. It was an absolute tendency, a property of the laws governing matter itself.

"The moment one tries to define what habit is, one is led to the fundamental properties of matter. The laws of Nature are nothing but the immutable habits which the different elementary sorts of matter follow in their actions and reactions upon each other. In the organic world, however, the habits are more variable than this. Even instincts vary from one individual to another of a kind; and are modified in the same individual, as we shall later see, to suit the exigencies of the case." (Principles of Psychology, chapter 4, p 104)

What differed for organic habits is that they can change. They are plastic, reshaping so as to adjust the organism to environmental variation. Anticipating the concept of neuroplasticity, James observed,

"Plasticity, then, in the wide sense of the word, means the possession of a structure weak enough to yield to an influence, but strong enough not to yield all at once. Each relatively stable phase of equilibrium in such a structure is marked by what we may call a new set of habits. Organic matter, especially nervous tissue, seems endowed with a very extraordinary degree of plasticity of this sort; so that we may without hesitation lay down as our first proposition the following, that the phenomena of habit in living beings are due to the plasticity of the organic materials of which their bodies are composed." (105)

For James, habits were the result of a process of channeling and pathing, in

which, like a trickle of water over time cutting a canal into the earth, expanding it until over time it runs deep, so do habits fix into place and determine a central tendency from which the organism settles its disposition. Such a habit might be seen as a variety framing ignorance, after all, we are often the victims of habit, unable to change and expand our possibilities once they have become fixed. At the same time habits entrench us, they also stabilize, allowing us to work within a subspace of the space of all possibilities with reliable outcomes.

Every choice logically entails an act of exclusion. The idea that cognition is eliminative is intuitive at first glance. Effective thinking rejects some points and leaves others. To be on guard against ignorance seeping into one's conceptions is to be critically minded. This is what critical thinking is: an intellectual power of rejection, skepticism and reevaluation. To be critical is to know that it would be ignorant to hold some beliefs without vetting them and to doubt incoming candidate beliefs until there is a good reason to believe them. Uncertainty is much like a vast body of water that surrounds a boat. The structural integrity of the boat must be continually maintained and any leaks patched, but there is always the risk that leaks may occur. Order must be kept. So cognitive agents tend to arrange their affairs to be reliable. Because coping with uncertainty is their main concern, and because uncertainty is nearly always present to greet them once they begin thinking about a problem, they must uncertainty-proof their affairs to the best of their ability.

So we sort and regulate things, and in doing so become sorted and regulated ourselves. The law mandates that automobiles be regularly inspected to avoid any future unpredictable behavior such as the brakes failing or engines giving out. Banks are regulated to keep them from behaving in ways that throw markets into turmoil. Job candidates are vetted rather than hired on the spot. We design things to protect against their behaving in ways we can't predict. The things on one's desk if disorganized can get lost any which way. Regularly sorted and regimented, entropy, unpredictability, is kept at bay. Cognition is about selecting a subspace of preferable possible configurations within the total space of possible states and then maintaining a commitment to that subspace of preferences to the exclusion of other alternatives. Without such entropic garbage collection and

bookkeeping, the agent's capacity to navigate its operating environment degrades. According to thermodynamic law the number of possible disorganized states of any system far exceeds the few workable, organized states. So affairs must be ordered not just to keep things from going wrong, but to exclude any deviation from the expected behavior.

Science fiction writer Philip K Dick coined the evocative word *kipple* to describe the useless junk and clutter that tends to automatically accumulate in the course of day to day activities. Laundry goes unfolded, dishes unwashed, unread emails pile up. Useless, disorganized stuff naturally accumulates unless managed. Such everyday entropic junk buildup is an invitation to uncertainty and more generally to increased *processing costs*. The more mess, the harder it is to know where everything is. Soon enough, the search space one has to operate in becomes harder to process as it ranges over all these extra haphazardly arranged options and irrelevant details. Efficiency means to keep things simple and tidy, to eliminate excess information so that only the relevant things remain. To work smarter rather than harder. It's about dumbing things down (simplifying) to such an extent it's easier to be smart.

It is not the kipple itself that is the cause of uncertainty. You can be certain there's a mess. The uncertainty arises from the increased complexity of the search problem while needing to find something in the mess. Just as there is material kipple, there is also kipple of the mind, confusions, red herrings, junk ideas, distractions, and just plain wrong assumptions that tend to accumulate the less intellectually tidy we are. Generally speaking, to think critically and arrive at truth is to cut through and filter out mental kipple.

Of course there are times when such protracted deliberative thinking does not occur and instead one has a flash of insight, seemingly skipping the eliminative steps to arrive at an idea without leaving behind a "paper trail" of strained reasoning. Such flashes of intuition, however, do not exclude the possibility that such eliminations did in fact occur unconsciously. Such eureka moments are expedited by having a clear headspace in which truths can surface unclouded or obstructed by diversions. Predictive cognition involves the serial ranking of various estimates, downgrading or upgrading them in conjunction to minimize

prediction error depending on how incoming sensory information matches sensory signals.

It would be foolish to simply wage an unending campaign against uncertainty for the sake certainty, to go blindly into the abyss without aim, constantly searching out and eliminating every last unclear detail about the world. The unknown writ large is not a monolith and cannot be encompassed as one. Humans are goal driven. They can only chip away at ignorance, but the question is where to do so. Given a set of goals, various unknowns are more vital than others. So rather than flail wildly in the face of the unknown and uncertain we create and stick to pockets of the world we can predict and fortify the perimeter around it that matter to us, only making occasional forays when needed into the surrounding landscape of mysteries and conquering new territory only where relevant or necessary. Ignorance is not usually just explored without direction, it is parsed into segments weighed by its relative utility.

The next chapter will examine the eliminative view of cognition through an analysis of trial and error behavior and artificial intelligence models such as supervised and unsupervised learning and the cognitive mechanisms of attention. These various perspectives on uncertainty-proofing can tell us much about how the mind, and mind-like algorithms, work in relation to ignorance.

Eliminative Cognition as Structural Ignorance

"But unfortunately we never grasp the import of any true statement until we have a clear notion of what the opposite untrue statement would be. The error is needed to set off the truth, much as a dark background is required for exhibiting the brightness of a picture." -William James

Trial and error process of elimination is the paradigmatic example of the aforementioned negative definition of knowledge acquisition. The known is what's left when all the other next-most-likely alternatives are excluded. The truth is arrived at through exhaustion and by the successive approximation through the serial binary separation of proof from disproof. Or rather, the

separation of the more likely to be right from the most likely wrong.

As in information theory each yes or no answer to each trial step reduces noise by some fraction. The process is similar to the game of 20 questions, in which an asker poses yes or no questions to an answerer and tries to guess the subject the answerer had in mind. By paring the problem down through binary pruning, the final value can be determined. An exponential amount of information can be eliminated with the most general questions, vastly simplifying the problem and constraining the range of topics addressed by the later questions. Essentially a well asked first question can cut the universe in half, and the next can cut that in half. Very soon you are left with just a few things to consider.

For the trial and error process of elimination/successive approximation to work, there has to be a prior hypothesis that extends over the entire set of initial, unfiltered observables. (As we will see later such a hypotheses assumes the maximum information entropy.) This hypothesis is general and inaccurate, but contains at least one true element within its scope, the needle in the haystack. One cannot arrive at a valid result by eliminating all, perhaps infinite possibilities. The process can only terminate at the correct result if the correct search space is first selected.

This selection problem beckons a vexatious haystack within haystacks regress problem. How to select which haystack out of all possible haystacks to begin the process? When somebody goes looking for their missing car keys, they do not start by searching in places they've never visited. Such a circumstance demands a hypothesis, which is indeed a positive, assertive proposition. You must have a vague idea where to begin. Elimination cannot proceed unless the vast bulk of possibilities are eliminated by this initial hypothecation step. Intuitively then, every investigation of a problem space must derive from some initial parameterization of it. A frame of reference or what Hintikka (2007) calls bracketing must be established to target the correct general search space. This is precisely the job of positivistic knowledge, to point the investigation in a general direction and tendency with a decent guess. Once this wide perimeter around the still hidden solution is drawn, positive knowledge runs dry, and the negative

elimination process takes over. Because if we knew any more we would have seized directly on the solution in the first place.

Such a hypothesis asserts only that the truth falls somewhere within certain boundaries. Without further eliminative work, the truth remains occluded by a fuzz of uncertainty within this bracketed space. Science proceeds by this tentative estimation process. It formulates hypotheses, runs a test, collects the latest data, cross-references it with past data, and either rejects, revises or maintains the standing hypothesis. Rarely is does science proceed through intuitive strokes of genius by through such a gradual process of approximation. With enough data and support and converging evidence a theory might be formulated. A scientific theory is then the residue of a long and continuous lineage of surviving hypotheses that managed to endure all previous eliminations. While nothing inductive can be taken with absolute certainty, we can at least bracket and zone out those null hypotheses which don't conform to the data.

Occam's razor is a close attendant to this bracketing process and in many ways embodies the eliminative principle. The more assumptions an explanation takes on, the more surface there is for errors, the less worth it has as an explanation compared to one that is equally supported with fewer assumptions. This very emphasis on explanatory compactness is a case in point of the eliminative simplification process that leads to reductive knowledge acquisition. In both the trial and error and Occam's razor cases, care is taken to exclude terms that are redundant, wrong, irrelevant or excessive. Simplification is then the epistemic act of leaving the least uncertain quantity for consideration, thereby unburdening memory constraints, and focusing attentional resources.

The old standby of mathematical and logical method, proof by contradiction, is another example of how to derive knowledge through negative exclusion. To prove by contradiction is to leave the proposition one aims to prove untouched. Instead one indirectly demonstrates how its logical opposite is impossible, self-contradictory or nonsensical, thus leaving no other alternative than the proposition one is trying to support. The demonstration does not so much provide supporting argument to the asserted proposition but instead leaves it as

a residue of logical elimination by disproving the counter-proposition. The meaning to be taken from proof by contradiction is that a proposition is not proven by direct evidence for it, but by evidence against the alternative.

Disproofs of a proposition's inversion act as a kind of protective buffer. Uncertainty, once again, only ever varies in its nonzero value, it is never zero outside of well-defined deductive formalisms. Thus, arguments for a proposition must be bolstered by arguments against its counter-propositions to both positively assert the claim and defend against its refutations. A perfectly supported argument does not just have positive evidence for it, but also has negative evidence for what's against it. Thus all sides of the proposition's logical space are fortified.

There are however some major problems with basing the truth of a proposition on arguing against its contraries, particularly regarding inductive inferences. (Deductive proofs by contradiction are generally final.) We may not know its counterarguments, we may not know how many of them there are, we may find the evidence for proving them false lacking. An observation statement ("there are two apples on this table") that can be decided by evidence has the benefit of direct empirical support. To what extent can we speak factually about a potential infinity of counterfactuals standing in opposition to any one statement of fact that is not immediately provable?

This is the problem which has been under constant consideration: there is uncountably many or perhaps infinite unknowns (this quantity is itself unknown and perhaps unknowable) relative to any finite quantity of knowledge. The problem of induction once again looms large. It is a problem of the incompleteness of information relative to a boundary of uncertainty that causes the inferrer to wonder whether or not they have acquired inquired adequate information.

NOTE: This paragraph between the carrots is likely to be cut and is on the chopping block .

^^^^^The doctrine of modal realism responds to this problem of infinite counterfactuals by asserting that there is a world where every statement is a fact —a logical multiverse in which everything that is false in one world is true in

another. A rigid designator like H20 picks out an existentially quantifies the set of worlds where this the referent exists. With regards to the contrast between knowns and unknowns, an empirical statement about an observation in this world poses a counterpoint to its inversion in another world. The logical content of the proposition leaves open its negation; knowledge of what is true in this world also provides knowledge of what is false in another. Positive knowledge, thus grounded by both its positive and negative consequences, provides a focal point around which an infinity of unknowns revolves, those unknowns belonging to the set of modal possibilities that must obtain in some world. So the more we know about this world, the more we know about not-this world. The unknown is contoured and shaped around the known. Any positive knowledge reorders the logical space of the unknown. It makes certain unknowns related to it inferable. ^^^^

The space of possibilities afforded by Lewis' plurality of worlds corresponds to the space of all possible unknowns. Lewis discusses this in his work on counterfactual dependence. In a world of uncertainty, modal reasoning is of essential value. Lewis, through his work on counterfactual dependence, establishes the similarity rules that must account for the most extensible quantifiers for locally similar possible worlds.*Counterfactual Dependence and Time's Arrow*(1979) Lewis gives a rubric for deciding the similarity between possible world contexts:

1. It is of the first importance to avoid big, widespread, diverse violations of law.
2. It is of the second importance to maximize the spatio-temporal region throughout which perfect match of particular fact prevails.
3. It is of the third importance to avoid even small, localized, simple violations of law.
4. It is of little or no importance to secure approximate similarity of particular fact, even in matters that concern us greatly.

In Lewis' conception of natural law, they are truths or collections of truths with the biggest "reach" over modal space. Natural laws are those truths which exclude the most counterfactuals. Universals, according to Lewis, identify

copresent properties across worlds. A giant red dragon does not share the world with us, but the property of *redness* does. Analogously, natural laws define the stable relations and causal orders across the maximum of possible worlds. More will be said about natural laws later, but in this case, it's informative to note that the the extent there is absolute certainty it is to be found in natural law. Every unknown contingent fact in a given world must conform to the collection of copresent truths that range over all its objects.

How exactly does knowledge reorder the relations to the unknown? Once a fact is ascertained to the highest degree of likelihood, the perspective on the world shifts. The modal space is rearranged, and certain counterfactuals are "deleted" or "taken off the table" as conditions in play. Facts form a network of dependencies. Any one fact must link to others, so knowledge of one fact shuffles its "sibling" facts into proximity. If it is known that Socrates is mortal, then it is immediately also known that Socrates has a date of birth, that there was a time before Socrates existed and a time after he existed, that there are situations (the ingestion of Hemlock, perhaps) that Socrates cannot survive, etc. In this way knowledge of a part leads to knowledge of a greater whole through chains of inference. In this way "seed" knowledge can provide a mechanism for traversing the space of the unknown.

Even if objective uncertainty presents an infinity of unknown possibilities, knowledge of general truths can exclude massive amounts of them in the same way that Lewis' take on natural law organizes the modal space. Knowledge acquisition is therefore much like the binary tree branching process mentioned previously where null branches are pruned and positive ones explored further. We do not have to waste time brute force searching for the facts that link to our known fact, those *relata* have been "brought closer" through the uncertainty reduction achieved through the acquisition of that fact.

Now having sketched this "outside-in" negative definition of knowledge acquisition, we next turn to examples in artificial intelligence to characterize the learning process further.

Use it to beat upon the cloud of darkness above you and to subdue all

distractions, consigning them to the cloud beneath you
- Cloud of Unknowing

To focus and attend to anything means to force out everything else. That's what attention is, an act of exclusion. Selective ignorance, *constructive ignorance* is therefore a necessary condition for effectual activity. "Paying attention" can be thought of as an act of shedding, of pruning a wide open decision tree until only the most goal relevant branches and stems remain. Attentiveness is the act of progressively diminishing the set of attention-grabbing stimuli until only the single prioritized residuum remains. Coming to attend to an object then involves a series of ranking operations in which one shifts between several candidate objects of attention, striking off those that do not fit the goal. In analogous fashion uncertainty as to what should be attended is reduced. Excess information, as a measure of disorder, is successively excluded at intervals in order to concentrate. In this way selectively *not* knowing is strictly beneficial to knowing. You must deliberately forget or become ignorant of everything else to focus on anything specifically.

Attention is critical to nearly every cognitive process. Attention can be considered one of the fundamental cognitive resources alongside memory. Because it is finite, how it is not used is as important as how it is used. It too is a negentropic process. It's the progressive "boiling down" of a diverse, chaotic information mass into a coherent, stable state.

When adequately attended the object of focus stands out from its background of competing stimuli. It is "cut out from" a surrounding ambient cluster of irrelevant objects and foregrounded. Thus to attend is to eliminate, and to allocate a remainder of content after excluding all that's irrelevant. Inappropriate details and noise is blocked out and the object of focus condenses out from a haze of uncertainty. Indeed, recent research into the neural mechanisms of attention have come to precisely this conclusion. The research indicates that the prefrontal cortex regulates the sensory thalamus (the part of the brain that relays incoming sensory signals to their proper destinations) through a specific basal ganglia pathway (Nakajima, Schmitt, Halassa, 2019). The details are unimportant, but the

discovery of this pathway which filters out distracting inputs ("goal directed background noise suppression" as the authors put it) lends strong support to the claim that attention involves imposing an ignorance structure around one's stream of consciousness in order to keep it on track.

To forget, to ignore, is sometimes to shed information purposefully to free up memory and make salient the matter of greatest import. Studies of neuro-atypical savants with perfect memory recall tend to note how this condition is debilitating. Their inability to discard memories and irrelevant details overloads their minds. It's not how much you know (remember), but that you know the right things. ADHD patients also illuminate the constructive role of ignoring by their deficit in it. A sufferer of ADHD cannot shed the competing objects of attention. They are instead vaguely aware of a nebulous diffusion of attentional objects, not one of which is ultimately selected as a focal point. They can't cut down on what they're selectively aware of so that only the most relevant priority remains. ADHD is in the last analysis, a failure of eliminative cognition. Indeed, neurobiological studies have generally concluded that the right hemisphere prefrontal cortex, which is involved in behavioral inhibition, is strongly implicated in the neuropathology of the condition (Arnsten, 2010). This suggests that ADHD is a disorder of the ability to stop, shut down, or cancel out active impulses, it is a paucity of constructive ignorance via decision tree pruning.

Such clinical examples lend support to proposal that ignorance can play a constructive role in intelligence. The pathology of ADHD demonstrates that attention is a negative process of exclusion rather than an effortful sustained positive concentration on an object. Or it is a combination of both mutually reinforcing positive and negative feedback systems. That concentration must be bolstered by the executive "willpower" (or the appropriate basal ganglion neural regulatory mechanism) that shuts out interfering stimuli.

It is in this respect that the limitation of the positivistic view of knowledge comes into view. The mind doesn't just run on what it knows, but what it doesn't know. Knowledge—or information- in the positivistic view is seen as fuel for the mind, its only source of content and the basis for its operations. But knowledge itself can only go so far. It needs the exclusion of competing hypotheses, a

negative inverse space, to balance itself against. Knowledge is contextualized by the negative space in which it stands in relief, a field of ignorance backgrounding a forefront of positive contents. This ignorance is not a blank canvas, but an organized substrate that plugs into knowledge through a complex interface. The modus operandi of scientific method attests to this fact. It succeeds because any scientific statement truth is pending further data. It remains "open" to the world of the unknown by the fact of its revocability and by doing so retains an effective epistemic stance. Given further data the statement can be revised in response to counter evidence, or supported with more evidence. Science therefore encodes the structure of ignorance through it receptiveness to its possibility and reality. Esoteric and perhaps even odd as this view of the relation between knowledge and ignorance may sound, over the course of developing the idea it will become almost intuitive.

what remains is the question: how do we come to terms with the limits of knowledge and the infinity of ignorance? How do we make use of uncertainty rather than let it rule our lives? And how do we make ignorance work for us rather than against us?

A possible objection to the perspective on the productive role of ignorance expounded here is that it is really about knowledge in a different guise, subtractive rather than positive knowledge. What we refer to when speaking about constructive ignorance is really the knowledge *not to do something*. This criticism is invalid for the following reasons. Knowledge and ignorance work in tandem in specific situations. To know what one should ignore is a subtractive act, which nevertheless leads to better process of knowing through omission. There are other examples of what McGoey (2019) calls strategic ignorance of the practical benefits in politics and in life of "playing dumb." Different allotments of information and the lack of information are distributed throughout the system to productive effect. This conscious commitment to ignorance is a distinct subject that will be investigated further. Suffice it to say that subtractive/eliminative knowledge operates in relation to ignorance. Many varied control systems and homeostatic processes throughout nature rely heavily on error-detection and

correction regulatory mechanisms which implicitly rely on this know-ignore subtractive relationship. The key here is that an error is known to be detected, the system actively adapts itself to the contours of what is wrong.

Secondly, the criticism that what we're really talking about in all cases is knowledge in different forms would be true if ignorance is essentially nothingness as in the common sense understanding of it. On the contrary here is a much stranger hypothesis, namely that of the *substantiality* of ignorance. If ignorance is in essence nothing, a negation of a positive term (knowledge), then Shannon's definition of information would be absurd. A nothing can't be reduced further. To speak of the reduction of uncertainty would be equivalent to trying to remove more from zero. Instead, the gain of Shannon information is like subtracting a negative number to approach less of a deficit in a quantity. Thus if the well-established theorems of information theory are to hold, ignorance-aliased to uncertainty - must have an ontological reality.

There is a weak and strong version to the claim about the objectivity of uncertainty/ignorance. The weak variant says that uncertainty is a psychological state, expressed through doubt and confusion. If these cognitive states are fully understood then the uncertainty is fully understood along with it. It is a state of mind, interesting and worthy of comprehension but not quite as profound as the stronger position. The strong version of the claim is that this uncertainty is manifest in nature by virtue of its complexity and fundamental properties. It is a property of the physical world and of thermodynamic entropy. We experience uncertainty because the world is not simple. From its complexity derives the concept of possibility, in that its thermodynamic properties affords a world where many alternatives can be conceived as viable.

In short we all wonder what will happen in the future. That's because we live in a world complex enough to prompt us to do so. This mystification about the future could have two causes. Either our cognitive state is deprived of the determinate facts—because that information doesn't exist yet or hasn't been made available— so we experience doubt. Or the future is inherently indeterministic based on a fundamental property of the universe. In other words, the present is the current edge of time, and the future from our perspective is

inexistent. Wondering about future events is then in essence to assume certain potentialities in the present. An uncertainty principle prevails in the universe, a fundamental limit to the powers of inference. If inference is bounded by iron-clad constraints, then we are left with an eternal struggle with uncertainty which defines the boundary between knowledge and the unknown.

Several popular techniques in AI are inspired by learning mechanisms believed to be employed by the human brain. One such technique is supervised learning in which the program is continuously corrected by an external "critic" until its output approximates the expected target in a manner reminiscent of behavioral reinforcement learning in animals. While effective for some problem domains, this method requires constant feedback from human trainers which places an upper bound on its speed, efficiency and autonomy. It also has limited discovery potential because the method only applies where the targeted result is already known.

In recent years unsupervised learning has taken precedence. This approach extracts patterns from the training data itself, enabling a more autonomous learning strategy with less data. This method has proven to be far more successful for problem domains in which the desired output is not known beforehand. Clearly this is a more powerful model of intelligence, as we generally suppose an intelligence to have a level of self-sufficiency and a capacity to produce novel conclusions for idiosyncratic problems. The success of the unsupervised method is instructive. It implements a mechanism in which past experience boosts the prediction of future experience and therefore reveals an interesting dynamic about knowledge acquisition.

What does it mean for knowledge acquisition that inferences about immediate inputs can be produced based on patterns extracted from the previous history of inputs without any external guidance? It builds a self-reinforcing structure of repeated adjustment cycles that fit incoming inputs to a repetoire of continually refined patterns. This would seem to support a positive take on knowledge acquisition. Indeed, there is little to argue against the idea that knowledge begets further knowledge, as per the discussion in the last section about the

ascertainment of a fact "drawing closer" its related facts and shaping the epistemic space into a more inferentially salient order by collapsing uncertainty. Prediction based on prior learning is indeed fundamental. The question is more about the surrounding edge cases— how does this process start?

Such "cognitive bootstrapping" as Windridge (2005) calls it, must begin with an absolute minimum seed value to subject to expansion. This seed value initialization itself cannot be derived through bootstrapping because the mechanism relies on the refinement of prior experience. Of course for the first experience there is no prior experience. In the beginning then is ignorance, pure and uniform, with no initial knowledge to gain purchase and bite into it.

In the case of *Homo sapiens* and other animals with higher brain functions, this seed value is provided by the genetic endowment. This extra-cortical information specifies the schematic template for the brain and provides it with the adaptive properties it needs to adjust itself to experience, as well as equips it with reflexes and other innate default functions to kickstart bootstrapping.

For unsupervised AI programs the weights of its parameters are initialized with random numbers. This is telling. It betokens this idea of the *potentiality* of randomness and disorder. Initializing the parameters with biased weights would over-constrain the degrees of freedom needed for the system to adjust to unexpected inputs. As the specificity of the initial parameters increases, flexibility decreases. Randomness taken to an extreme in a sense is a catch-all, it contains every possibility each of which however cannot be utilized as they have not been extracted. In somewhat similar fashion, neurologist Karl Spencer Lashley's theory of the equipotentiality of neural tissue claims that any part of the brain can potentially learn functions attributed to other functional areas in the event that that area is damaged. Clinical studies of substantial recoveries of function in cases of severe brain trauma lend credence to this hypothesis (Nelson et al., 2016). Over-determined structure is detrimental to adaptiveness.

The fact that the world's causal structure exhibits regularities according to the laws of physics is the grounding principle of all pattern detection algorithms in the obvious sense that it is what puts the patterns there to detect. The problem is how to tell a sequence of events really is a pattern. Cognitive systems derive their

power to generate inferences from such regularities, but only by contrasting it and extracting it from informational kipple and excess. A pure chaos is unpredictable in its entirety, while perfect order is undifferentiated and uninteresting. Some degree of chaos is the prerequisite to the evolution of intelligence because a perfectly predictable world does not need intelligence to navigate it. Intelligence is invoked in a state that straddles both extremes of order and chaos. Such a world is both organized and dynamic enough to both permit and prompt problem-solving innovations.

It's this very thermodynamic middle ground that establishes the necessary and sufficient conditions for the evolution of predictive inference algorithms bounded by an atmosphere of uncertainty. Only that there is uncertainty, therefore, is the need intelligence and inference in the first place. Both perfect order and perfect chaos are uniform in their simplicity. When combined and counterbalanced in the right ratio, the conditions of possibility for knowledge arise. Intelligence is then, both a reaction to and an instance of complexity. The possibility of being wrong is the only thing that makes being right bear any meaning. When one can't be wrong, when the problem in other words is maximally trivial, the inference is valueless and there is no need to be smart. Or rather there is no need for inference machinery at all, and thus there would be no inference. For precisely this reason uncertainty and ignorance are the form-giving prior conditions for intelligence. Information entropy is the raison d'etre for thinking.

The very notion of thought occurring presumes that a pocket of ignorance and uncertainty must be collapsed. Thought "drives into" ignorance, struggles to overcome it and is "pulled forward" by it in the process of eliminating it. Tasks that do not require thinking are reflexive and subserved completely by programmatic instructions or static, unadaptive mechanical routines . Bridging the gap between knowledge and ignorance is inquiry, which is what allows the former to gain purchase on the latter. While this is not yet the place to discuss the logic of inquiry, it's important be reminded that it is waiting in the offings and is a crucial piece of the puzzle.

There is a potential for *richness* in randomness. Effective cognition is "anti-

fragile" in the sense that it gains from disorder (Taleb, 2012). Cognitive systems "surf uncertainty" as Clark puts it, at times deliberately seeking out the unexpected to calibrate their predictive estimations. Such systems turn their "uncertainties and expectations into concrete objects apt for test and challenge." ([https://aeon.co/essays/use-uncertainty-to-leverage-the-power-of-your-predictive-brain]). To turn uncertainties into "concrete objects" is quite the revealing phrase. It recalls what we've been alluding to this entire time, namely, the definite reality of negative epistemological contents such as ignorance and uncertainty.

Randomness only becomes truly detrimental when it is *volatile*: when how things change itself changes. Pure instability has little to no discoverable structure, and thus the conditions for securing accurate predictive estimations aren't satisfied.

Generally speaking, the more a cognitive agent opens themself up to the vast variety of possibly relevant experiences adjacent to its regular routines, the more dynamic, flexible, and responsive it grows to sudden spikes of unplanned contingency. We usually "test the waters" before jumping into the unknown. A cognitive agent that is rigidly habituated to a single niche becomes locked into that routine, and therefore suffers degradations in adaptiveness. To break with habit completely, however, is also potentially harmful. Habits exist for sensible biological reasons: the existence of a habit itself is evidence that the course of action describing it has been proven to be a viable survival strategy in the past, simply by virtue of the fact that the habit has not resulted in death or injury. In the realm of familiarity the agent most navigates most frictionlessly, but too much comfort can undermine the ability to respond to change. Too much overextension, however, "being out of your element" or a "fish out of water" is equally disorienting. An intelligence must therefore thread the needle between absolute spontaneity and absolute regimentation, using the former to augment and extend the latter and to therefore build its capabilities while maintaining a reliable base of predictability.

Thus far the topic for discussion has been almost entirely dedicated to the nature of uncertainty. What, however, is certainty? Intuitively, the argument for

the objective existence of uncertainty would seem to also imply the existence of certainty, in the same way *off* is defined only in relation to *on*. The next chapter will take a different direction entirely. It will examine Wittgenstein's tract *On Certainty* and how this fits in with the overall picture.

Part II, Section I: Physical and Computational Constraints on Knowability

This has brought us to a critical juncture. The next two chapters on physical and computational (also logical) limits are crucial for defining the absolute dimensions of ignorance. If there are indeed physical and computational limits on what can be known, these upper boundaries serve as a first approximation for characterizing the constraints on knowledge.

It must be stated at the outset that although it is still by far no small task it is much more plausible to define the limits of what is knowable than the unknowable. At which point, the hope might then be to try to infer the unknowable from the knowable. This is because what is knowable can always be identified, whereas what is unknowable—by tautology—cannot be known and therefore can only be inferred indirectly. Perhaps all that is knowable share one essential property that makes them knowable, and thus we can conjecture that anything which lacks this property is unknowable. At minimum the knowledge limit has a surer possibility of being determined, because we can at least establish that what is inferable is bounded and finite based on Wolpert's proofs. The same cannot be said about the unknowable, as its scope is undefined. Given the logical asymmetry of knowledge and ignorance, we cannot necessarily infer that it too must then be bounded and finite, given that ignorance, according to the premises, is not merely potential knowledge.

(In what follows I will assume that knowability and unknowability are intuitively amenable to a set-theoretic analysis, that they can be represented as collections of units organized by some shared class membership property. It's not clear that it makes perfect sense that either knowns or unknowns can be

quantized into such atomic propositions, and especially so for the unknown, this, but it certainly aids clarifying thought about the subject.)

The knowability versus unknowability problems have distinctive forms. After all, everything we know is knowable, all known items are immediately included in the set of all knowables. This feature imparts a certain salience to that class. What is unknowable is conditioned on what is knowable, what is knowable speaks for itself. Nothing that is knowably unknown can be derived without knowledge. Inversely and obviously, not everything that is unknown is unknowable. The two classes, knowable and unknown therefore have distinctive acceptance criteria. . Further, nothing that is unknown can be directly inspected and sorted into its respective bin without immediately transforming into knowledge or being related to knowledge by the act. As soon as we become aware of an unknown, it becomes tethered to that fact; it must share come correspondence with our knowledge and is more likely to be in principle knowable. Consequently the set of unknowables is drastically harder to build. In fact even in the course of searching for unknowables we continually add more elements to the set of knowables as a side effect, trying to discover what is unknowable, a seeming paradox, automatically gives knowledge.

In terms of knowability, the crux of the issue is that no matter how many pieces of factual knowledge are to be enumerated, it cannot be inferred from this data how many are elements are remaining in some conceivable set of all possible factual knowledge. (In some way this is at the root of the problem of induction, its impossible to tell you know enough without knowing you've exhausted every possibility.) No matter how much knowledge is accumulated it is impossible recognize the difference remaining between incomplete knowledge and complete knowledge. We can't even deduce how much ignorance is remaining from this act, as if one could know how much volume in a cup remains to be filled to the brim with water by comparing the difference between the fluid level and the height of the cup. Ignorance is bottomless in this way. No matter how much knowledge is gained, it is impossible to specify a correct ratio between what is known and unknown from it. This holds true even if we consider a single individual trying to learn what is already known by others, for

example a student studying an established field of knowledge. Because within that field, the borders are constantly being extended through discovery and new findings.

So even if it were remotely humanly or physically possible to approach that limit of knowledge, the edge case on the other side could never be reached. (As will soon be shown the issue is even more confounding when it is considered whether the set of all facts is countable or uncountable.) And how much knowledge we have certainly cannot tell us whether the set of all factual knowledge is of equal size and has all the same contents as the set of *all* facts, including unknowable ones, should they exist. If it is a subset, then there are unknowable unknowns, if not, there aren't. However, it is itself unknowable, or more precisely undecidable, whether the set of all factual propositions and the set of all facts have the same cardinality for the reasons just stated in the bottomless cup analogy. As we shall see this one undecidable proposition, and others, appear to be the only types of elements from which the set of unknowables can be reliably built.

It seems to be easily overlooked that there can be independently factual propositions *about* propositions regardless of the factuality of their "inner" proposition it references. The secondary "outer" proposition's assertion can be decided true or false even if the first proposition cannot. A key example of this would be Gödel's numbering method which allowed him to encode true, decidable statements about what is undecidable in arithmetic. Such secondary, metalogical or metamathematical propositions can therefore be included alongside or set aside from the first-order facts to which they refer and which are not evaluable (incomplete). There are in other words factual propositions with non-propositional referents and factual propositions about such propositions. An excellent example of this situation is the above statement about whether there is parity between all possible factual knowledge and all facts . This has an undecidable truth value which a second, decidable proposition can assert, namely that it is impossible to confirm that one has come to know everything. But by posing this very question, it tips the scales. Gödel's theorems constituted exactly such a set of such metalogical statements.

There is also the problem of how to sort facts of mixed epistemic status. Is a known unknown to be sorted into the bin of unknowns or knowns? In some sense does it not belong to both? Or must there be only pure knowledge that is counted, known knowns? The answer to this question is not at all obvious.

Such propositions are arguably undecidable, they cannot be answered. The proposition can be known and formulated, but its answer cannot. As a rule every undecidable proposition means that this is a yes or no answer which cannot be known. these meta-propositions about how to classify the knowable versus the unknowable are subject to the same rules as the first-order propositions about what is knowable and unknowable they seek to sort out. The seeming solution is to separate both propositions and place each one in the respective bin, but the bin that contains the second-order proposition must always reference the first order proposition, the logical sign of which is always the opposite of what members of that class should have. So the the set of what is knowable is built out of first-order factive propositions and second-order propositions that introspectively evaluate the knowably unknown. But just because it cannot be known that they are undecidable, does not mean they are in principle undecidable. They are really *undecidably decidable*, a compound term which bears a logical distinction from a merely undecidable or decidable proposition. In other words, the proposition is unprovably decidable. It can neither be stated that the inner proposition is knowable or that the outer proposition that wraps it can know it can be known. It should also be stated that working out the meaning of this riddle resolves, or more humbly at least sheds new light on Fitch's notorious paradox, which in short asserts that if all truths are knowable all truths must be known.

The most effective method to the problem unknowability is therefore to analyze it in terms of constraints on knowledge, which are themselves knowable facts within the set of knowable knowns that define what rules it must follow. Better to question the how the generation of knowledge is restricted because this promises to be a more compact problem. Anything that disobeys those rules must be unknowable. This allows us to neatly stay within what is knowable and still make confident assertions about what is unknowable. If constraints on

knowledge are defined, we do not have to, impossibly, enumerate all that is knowable to know what is knowable and therefore what is not. Because we can go by the rule for what is knowable rather than iterate through all facts, which again is impossible.

First, there is a tendency to assume that questions of the limits of knowledge can be answered on strictly logical terms, as if the physical universe, as if empirical reality, was of no concern. What then is logical truth if not an eternal tautology expressing nothing of any value? How can truth be separated from causality? Logic imposes on itself its own criteria for truth circularly, when it is itself an illogical proposition that the limits of the mind must reflect the limits of logic. This "mirror of nature" rationalist assumption that logic reflects the rational order of nature is not a tenet that can be asserted within logic as a necessary truth. Hence so many strange edge cases in epistemic logic such as logical omniscience that seem to make it fall short of its aims of formulating a rigorously accurate logical model of knowledge. Is knowledge we can never have really knowledge even though it is logically necessitated? Conversely, what about facts of the physical world that do not fit logical rules? Quite the contrary, the ultimate limits of logic are tied to the the ultimate limits of the mind which are tied to the ultimate limits of the brain which are tied to the ultimate limits of computation which are tied to the ultimate limits of physics. So it's beneficial to look at what rules of *physics*, which is an ontological prior to such computational, neurological, cognitive and logical rules, has to say about constraints on information access which must also speak to constraints on knowability inasmuch as one must be informed to know. Secondly, it's beneficial to look at computational limits, as this provides inroads into the rest of the preceding terms that physics cannot while providing sufficient means of generalizability. All in all though it is the section on computational limits in which it is clarified by logical analysis which I believe is more fundamental and directly applicable to the problem of unknowability. That's because one cannot derive a concept of knowledge or its modal variations within a discourse in terms of physics even though it depends on it causally.

The Physics of the Unknown

Physicist John Wheeler once said his career could be summed up in three phases. First everything is particles. Second, everything is fields. Third, everything is information. This "informational turn" in physics can be taken to mean that the physical universe is essentially constituted as an epistemic construct. There is a continuity between things/process/events such as minds and those such as black holes. Because this so called It from Bit is such a going concern in cutting edge physics, it's justified to see it as direct support for the argument that uncertainty and ignorance is a property of objective nature and that a structural incompleteness of information is inherent to the subject-object relation which defines the epistemic standpoint of an (un)knower to the world.

The laws of physics are unique and powerful achievements. To the extent that they describe the rules by which everything in the universe must conform they also express the upper bounds of empirical generalizability. The laws of physics tell us how any one generic class of phenomena will behave given such and such conditions, but they cannot tell us how a particular object will behave other than with respect to this very general rule structure. In this way the laws of physics are both powerful expressions of positive knowledge and of the limits of such knowledge. They specify a framework in which the possibilities of natural phenomena are constrained, but they do not speak to any individual instance of a physical event within that framework. Unless, of course, we go out of our way to identify a particular object and calculate how the laws influence it.

Given the fundamentality of physics it is philosophically appealing to try to ground many concepts in physicalist explanations. The question at the forefront for our purposes is if *epistemology* can be so grounded. The answer I propose is yes. Modern physics is thoroughly *informationalized*, that is, drawn up in the mold of Wheeler's It from Bit view. Information that is to say is popularly considered among working physicists ontological to modern physics: not merely a descriptive property of systems but of the essence of physical phenomena as such (Beraldo-de-Araújo, Baravalle, 2017; Lloyd, 2006). Information is central to the understanding of modern physics, and it just so happens that it is also the

governing concept of modern epistemology too. As Hintikka argues, the question of knowledge acquisition can be substituted for information acquisition to the benefit of greater conceptual clarity (Hintikka, 2007). Information is quantifiable in a way that knowledge isn't and its definition is precise. Furthermore although you can have information without knowledge you cannot have knowledge without information if the brain is an information processor and knowledge is a property of the brain. Computers don't have knowledge, but work with information, the same sort of information that is the main ingredient to knowledge.

There are sure to be several key physical concepts that are relevant to explaining a naturalistic epistemology. The one of primary importance, however, is entropy.

The Second Law of Thermodynamics states that the net entropy (disorder) of any closed physical system always increases over time. Unless supplied with a countervailing energy source, the sum of free energy available to do work within the system will tend to dissipate into an unusable form. In short, everything runs down. Heat, expressed as the rate of motion of molecular ensembles, always tends to level off toward a "universal room temperature" of thermal equilibrium. This happens because every particle within the system transfers some of its energy to those adjacent to it, which in turn does the same, until the total energy of the system dissipates into space and a homogenous uniformity as the energy of the system becomes more uniformly distributed.

Given the second law's profound relevance for many natural phenomena and its applicability to many practical domains in science and engineering, it is natural to wonder how entropy applies to the brain as the organ of thought and behavior. Quite obviously the brain is a physical system, using energy and producing heat and thus subject to entropy. Yet is this connection meaningful in the sense that the brain's relationship to heat dissipation has a role to play in how the mind works?

While the brain behaves as a typical energetic system it is less obvious how its functions might be explained by the same principles. The internal dynamics of

the brain means that it can self-organize by transferring energies from one part to another through the transmission of electrochemical action potentials. As energy flows through the brain more structure is imparted to its neural net and its computational options expand as it thereby learns. Yet the brain is unique in many respects as an energetic system. It is a representational engine. As such it "looks" at entropy while at the same time riding the waves of entropy as it transitions from usable to unusable form. It comes as no surprise that entropy as the measure of information (descriptive complexity) needed to define a system's state would be relevant to an information processor like a brain. Hence the view that the brain is essentially an entropy reducing (negentropic) machine. Indeed research that has looked into the relationship between the brain and entropy suggests that the same principles that guide the accumulation of self-organizing structure in other domains may apply to the brain as well. The question becomes, how the brain as a physical system can change to accomodate new structure. As some research suggests, the role entropy might play in brain development is akin to a "noise injector" that adds a degree of nonlinear behavior into the stable state configuration, allowing for the possibility of new configurations. As Stevens, et al write,

> The overwhelming influence of entropy manifests itself as a critical instability. As the system approaches a critical instability, constraints dissolve thereby freeing previously cohering parts and causing them to interact. These interactions may allow the system to reach a new configuration, better suited to disperse entropy. Therefore, in self-organization, new structures are predicted by an increase in entropy (indicative of critical instability) followed by a decrease in entropy (indicative of the onset of new constraints). In nonlinear dynamics, this sequence of changes is a phase transition. (P 1812 https://www.researchgate.net/profile/Damian-Kelty-Stephen/publication/40455355_Dynamics_of_Representational_Change_Entropy_Action_and_Cognition/links/0912f50bccf485bc58000000/Dynamics-of-Representational-Change-Entropy-Action-and-Cognition.pdf)

This "constraint breaking" proves that entropy is not an entirely destructive tendency and that the increase in disorder within a system allows it to access other regions in its "phase space" of potential states by disrupting its static routines. This in consistent with what was said earlier that uncertainty can be a

resource, that by injecting a bit of randomness into a process it may behave more interestingly than if it settles on a stable, repetitious state.

The riddle then becomes how entropy in the brain leads to a subsequent gain in order and representational complexity. This has to do with the relationships between physical and information entropy. Essentially, by acquiring more self-organized structure with the help of entropy expanding accessible states within its phase space, the brain gains greater ability to represent more complexity in the world. It acquires more power by its gain in structure. Information entropy is the correlate of uncertainty; the probability of an event is directly proportional the amount of surprise, or unusualness of a data set. High levels of disorder increases the computational search cost of finding the fact and as such, the more complex a context, the less rationally certain we are of the facts, the more thinking (entropy reduction) must be done. The very tension created here between the difference between the entropic complication of the computational search space and the processing power needed to extract the signal from the noise is what provides the evolutionary rationale for advanced intelligence. The increase of disorder can provide for a local increase in order for a far from equilibrium thermodynamic system, of which lifeforms are seen as a striking example (Prigogine and Stengers, 1997 (https://books.google.com/books?id=-VI8093PJuUC). Life are thermally open systems, they exchange matter and energy with their surroundings, and as such as energy flows through them and dissipates life "rides the wave" and profits off the energy flow as it goes from a more usable to an unusable form. As energy flows through such a "dissipative structure" a local gain in organization can be had even though the net tendency of the entire physical system (organism plus environment) tends toward equilibrium. To translate this into information terms, the introduction of a degree of randomness and uncertainty (ignorance) is actually conducive to learning and discovery and without it the brain would tend toward stasis.

Now to turn to a more definite description of the links between information entropy and intelligence and expand on it in more detail.

As stated previously role of entropy in thinking more generally is not to be

underestimated. Thinking has a definite physio-informational correlate in entropy. Information entropy measures the degree of surprise value (surprisal) in low probability versus high probability events. The probability mass function specifies that the lower the probability of an event the higher its information value. Meanwhile redundant, entropic high probability events have less value. There is something intuitive about this formulation. We rarely take note of the "same old, same old" of our everyday lives, but as soon as something different happens like a wedding or childbirth, it's remarkable and newsworthy.

The detection of a low probability event corresponds to the extraction of a signal from the signal to noise ratio. Certain methods are used in communications engineering to dampen noise and isolate the signal, such as lossless compression in which more probable, that is redundant, data is compressed into an expandable form. For example the sequence xyz-xyz-xyz-yz-yz-z can be compressed to (3)xyz-(2)yz-z. This both reduces the chance that information will be lost during transmission and makes the data take up less memory space. Does something similar happening in the brain?

It is well known that the brain becomes habituated to redundant stimuli. The more common an occurrence, the weaker the response. Indeed, the higher the chance that the stimulus will be ignored. Change blindness is the phenomenon where a person fails to notice a change in a setting because the change is so subtle that it does not trigger surprisal (Simons and Levin, 1997). In visual perception, much of what appears in the visual field is not actually present in the stimulus, but is "cached" in the visual cortex and filled in by feedforward predictive signals representing what the brain expects to see. Except for the immediately attended to focal point, peripheral vision is patched in with repeated, low resolution visual units. These features make the brain susceptible to change blindness as unattended components of the visual field "carry over" from the prior state information held in memory. Because nothing important seems to have changed, the brain merely reuses what it had already rather than invest energy reconstructing the scene afresh each moment.

Thus what we actually perceive does not correspond to the immediate visual scene but lags behind it until sufficient surprisal triggers an update signal by way

of prediction error. Error, in this case, is how the brain readjusts and corrects its model. While the operation of the visual system might not constitute compression in the same sense as in machine computing, it does represent the brain's encoding of the statistical redundancy of features present in the visual field. Such redundancy encoding is a constraint, in which only those uncertain, surprise laden elements are targeted while the most common features are automated away from direct conscious processing.

A common theme of cognition is the sustainment of an invariant holding pattern of redundant expected values and the subsequent reconciliation of punctuated, unexpected values that deviates from the held pattern. When we perceive new information, when we *notice something*, what we detect is the divergence of an occurrent value from the expected value. We see the relation between the invariant background and the element which "goes against the flow" and thus learn something new. Feedforward prediction signals are contradicted by feedback signals. Learning therefore is a form of contradiction; the anticipated pattern must be rebuked for new information to be recognized.

Uncertainty can then be seen not simply as what makes reasoning and decision making challenging. It's also what fuels learning. Without atypical signals the brain is won't to recycle its already cached representations without updating them in the face of novelty. The cognitive decline of senility can be seen in these terms. With senility cognitive flexibility degrades. A lifetime of accumulating expected patterns ultimately "maxes out" and the number of stimuli signals that trigger surprisal diminishes (along with an associated decline in performance of the nervous system). Learning in youth in contrast is accelerated for the opposite reasons. There is more novelty and there are always more surprises. Learning therefore corresponds to the amount of info-entropy that is perceivable to the brain.

More generally the entire computational schema of domain general cognition (general intelligence) is designed around entropy processing. If there is any universal law of cognition it is to sort the wheat from the chaff, to discriminate goal relevant sensory features from the sea of noise. The very act of cognitive deliberation is to cycle through hypothetical possibilities and successively strip

away implausible alternatives. The function of the brain is to bestow order on an initial, fuzzy chaos. In early childhood the world presents as a "buzzing, blooming confusion" as James put it, only for time and experience to sculpt that disorder into a coherent, navigable scheme.

Indeed recent research seeking to connect fundamental physics to consciousness interprets the latter as the maximization of available information in the brain (Erra, Mateos, Wennberg & Velazquez, 2017). In this view, consciousness is directly associated with the complexity of the brain. As the authors put it, "normal wakeful states are characterized by the greatest number of possible configurations of interactions between brain networks, representing highest entropy values." It would seem at first glance that the view that consciousness represents the highest entropy state confutes the view that cognition is entropy reduction. Yet consciousness and cognition are far from synonyms. In actuality this aligns with what I am proposing. Namely, that consciousness is possibility-oriented and possibility maps to higher entropy. Most knowledge and thinking is in fact not conscious, as the discovery of the cognitive unconscious shows. Consciousnesses' function is to "cast the widest net", to make a state of ambient possibilities available so that some particular stimulus of surprisal interest is captured and attended. Ambient, Default Mode Network (DMN) consciousness is a representation of the highest degree of information exchange in the brain, the state which contains the most complexity, openness and possibilities. Such a default state is optimal when you are facing a world full of the unexpected and must be sensitive to any number of eventualities.

The brain's use of redundancy in structuring its perception of the world is ubiquitous. We have just seen how low grade processing in the visual cortex uses statistical redundancies to fill in for the repetitive features of the visual scene. On a deeper level redundancy shows up in the most rarified of mental abstractions.

Consider analogical reasoning. Researchers in the field of cognitive analogy go so far as to call it the "core of cognition" (Hofstadter, 2001). At its root analogy is the mind's way of reusing patterns in order understand what it doesn't know in terms of what it does know. This ability to leverage existing representations to

accomodate more challenging unfamiliar concepts is powerful. Analogy detects similarities in the organizational structure of both the source and target representations. Such "structure mapping" is no mere figure of speech, but a central device of intelligence. Such mappings (like entropy to thinking or of information availability to energy flow, for example) are the fundamental forms of abstraction for much reasoning and theory.

By analogy the mind also reduces uncertainty by fixing on invariant representations and recycling analogous schemas across the compared instances. Through such mappings of parallel structures, the cognitive operation of analogy economizes its representational space by *folding it onto itself* in a manner of speaking. That is, by making good use of its contents to knit together the broadest shared patterns between different representations. Cognition, in computer science terms, *minifies* its representations by extrapolating invariant forms across as many comparators as possible.

Mathematics perhaps of all represents the idea of redundancy the most. A major goal of mathematics is to discover the most compact, general forms for classes of mathematical objects. Mathematical theorems are a master class in the economy of thought, perfectly idealized generalizations. Each theory in a sense contributes to the structuring of mathematical space; each theorem reorders the whole of mathematics and makes further reasoning about it more frictionless, clearer, and more expedient.

The important thing is not what is known, but that remains unknown relative to it. knowledge in one sense is stasis, low entropy order. It is the point at which action ceases and the schema persists across as many generalized objects that it can support until "disrupted" by surprisal. To the extent that conscious awareness is a maximal entropy state, knowledge is a minimal entropy state. knowledge acquisition specifies a folding up of the representational space, a reduction in possibilities. Therefore by knowing something we can think less about it, exploratory effort is saved. All the convoluted effort to learn and establish the concept is discarded and a straight line to the point achieved. Conversely so long as something is not known, thinking must happen; entropy predominates. There is too much murky possibility clouding our ignorance that

must be sorted out.

Having now characterized the premise that entropy and cognition are deeply related, the next step is make this connection explicit.

In 1957 physicist E.T Jaynes published a pair of papers that presented a provocative and novel interpretation of statistical mechanics. These papers recast it in terms of information theory. It connected the use of statistical inference used in the physics of macroscopic ensembles to the concepts of Shannon entropy and Bayesian constraints on inference under conditions of partial information and incomplete prior knowledge. Jaynes introduced the concept of the principle of maximum entropy as a means of quantifying probability in terms of the absolute measure of uncertainty inherent in a system. Jaynes' ultimate achievement was to peel off the logical schema that was fused to the statistical physics described by Gibbs and generalize it as a form of logical inference applicable to an open-ended variety of domains as paradigm of probabilistic reasoning. As he put it, "In the problem of prediction, the maximization of entropy is not the application of a law of physics, but merely a method of reasoning that ensures no arbitrary unconscious assumptions have been introduced." (Jaynes, 1957 I, p 630).

In effect the principle of maximum entropy regards the initial, or canonical, distribution of an ensemble (a set of particles or more generally any rule-defined logical element capable of undergoing transformations) before anything is known about as the state in which all likelihoods are equally probable. This canonical distribution is associated with but not the same as Laplace's principle of indifference, which states that in the absence of any means to differentiate the variable likelihoods of a given set of outcomes, rationally all outcomes should be afforded equal probability. This is the same as saying that we are at the most ignorant when anything goes, when the selection of a description of the present and future state of the system is arbitrary. So best to play it safe and assume anything can happen with equal odds, this would be less presumptuous than to assume drastically divergent odds randomly or by flimsily contrived guesswork. Measuring the maximum entropy of a data set gives an objective characterization of the amount of uncertainty contained in the system. This means that the

maximum entropy measures the maximum degree of possibilities represented by the *configuration space* of the system. It captures any possible state that is derivable from all the degrees of freedom, all the ways each element of the system can swing within the system based on specified parameters. Parameters such as, to use a simplistic example, any one of a number of point particles being able to move up, down, left, or right on an idealized 2D plane.

Jaynes devised a sophisticated mathematical framework to explicate these ideas which he developed over the course of his career. Prior to his death he was working on a massive tome on probability called *Probability Theory: The Logic of Science*, which, although he unfortunately passed way before completing it, still clocks in at a staggering 750+ pages of detailed mathematical exposition. It would be dishonest for me to claim to possess enough background knowledge to give a full picture of these highly elaborated and technical formulations. However his key findings and specifically his broader interpretations of them are profoundly relevant to the conception of ignorance I am propounding, even if they do not perfectly align with it.

All told, Jaynes was a committed determinist and a bit of a philosophy-phobe. He refused to believe that probabilistic states could be attributed to the physical system proper and instead insisted they were a reflection of our own ignorance. According to him the quantum uncertainty principle is merely a measure of our incapacity to penetrate into the hidden physical causes which prevent both position and velocity from being measured (observed, known) simultaneously. In this respect he was in line with the subjectivist school of probability, popularly represented by Bayesians, who view probability as a consequence of incomplete knowledge and inferences derived from prior belief states, versus the objectivist school, supported by frequentists who assert that the relative frequencies of the different outcomes of a distribution is an objective feature of how patterns of events unfold in time. Subjectivists view probability as as essentially mental in nature, quantifying the expectation values of belief. Objectivists view the likelihood that an event will occur as a property of the observed behavior of the events over a history.

The views espoused here align more with the objectivist school, but only in so

far as I argue the counterintuitive claim that the conventional view of ignorance as a mind-bound property is false and it in fact exists objectively. I do not see the anticipation of states based on subjective expectation values seen in the subjectivist school as taking a sharp departure from the calculation of frequencies in the objectivist school. Whether you say this means the subjective subsumes the objective school or vice versa is meaningless. Jaynes' was against the tendency to think that the use of the various probabilistic mathematical constructs in science meant the observed phenomena themselves were inherently probabilistic in nature. He was against the tendency to view "random variables" as designating actual phenomena rather than as representing a label of our own ambiguity. To confuse, in other words, one's own uncertainty about the system with inbuilt uncertainty within the state of the system. To think so was an instance of what he dubbed the *mind-projection fallacy*, the tendency to erroneously mix up how the mind works with how the world works. I respond elsewhere to this potential criticism in the anticipated criticisms section. Suffice it to say that this fallacy only applies if there is a hard and fast binding of ignorance to mental states. But if a physical system or data set in general at maximum entropy would is equivalent to any possible observer's state of maximum ignorance, this implies that their ignorance is an objective function of the data rather than the state of their minds.* That ignorance supervenes on entropy*. They may not be identical but they covary in strict parallelism. Why else would an utterly simplistic data set carry so little uncertainty? Why is the complexity and heterogeneousness of the data a proxy for how uncertain it is? Uncertainty is a function of physical complexity in that any cognitive agent exposed to it would more likely be put an uncertain state than if the system contained only minimal possibilities. This is obvious on the face of it, the more possible states a system can adopt, the more there is to be ignorant of. Therefore I view the subjectivist/objectivist dichotomy to be false, a mistake that stems from a misconception about what ignorance is.

If probability is a measurement of ignorance it's natural to wonder what it is actually measuring. Certainly this rules out the notion that ignorance is nothing. Measuring the lack of knowledge is senseless, for one thing we don't even know how to measure knowledge. Nothingness or absence can't be quantified without

comparison to a definite positive point of comparison and a definite limit. For example, one can say a glass of water is 2/3rds full and therefore 1/3rd empty, using the positive quantity of the water as the point of comparison and the explicit dimensions of the glass as the limit. knowledge however has no definable upper limit. You can't say you have 2/3rds of the knowledge, as that supposes you already know how much knowledge there is to have. To know how much there is left to learn, however, you would not be ignorant and in need of learning to begin with. So unless ignorance is somehow embedded in the structure of reality quantifying the ratio between knowledge and ignorance is not only absurd but impossible.

Does information come to the rescue? Information has the benefit of being eminently quantifiable. However, as will be stressed repeatedly, information is not knowledge. At most information is the ingredients of knowledge, the materials which it correlates, assembles and associates in a structured representation. Thoughts may feed on information, they may in some mysterious sense be made out of information, but they certainly aren't information. "You are what you eat," the saying goes, but the catch is of course that the amino acids and vitamins of the consumed food are subjected to a transformation function in the process of organic assimilation in digestion and reorganized into the pattern of a qualitatively different thing. As Brillouin, one of the first to explore the deep connections between information and entropy wrote, "we may state that information is the raw materials and consists of a mere collection of data, while knowledge supposes a certain amount of thinking and discussion organizing the data by comparison and classification. (Brillouin, [http://www.physics.mcgill.ca/\~delrio/courses/phys559/lectures%20and%20notes/Brillouin-Information-Theory-Ch1-4.pdf], 1957, p ix). Food does not eat food; neither does information think information.

If ignorance does indeed tend to be increased in the shadow of increased complexity, One can still be a Bayesian and maintain that the complexity and irregularity of the data set is what objectively necessitates its mechanisms of probability estimation. An utterly uniform and simple data set does not take much time for a Bayesian inference engine to figure out. In *Probability Theory*

Jaynes takes particular aim at some of the well known faults of Bayesian algorithms, namely what might be called the initialization problem: how to select the right initial priors without regressively assuming priors for those and without having a preexisting belief model of the data; and what might be called the bootstrapping problem: how to confirm that the selected priors are an appropriate fit for the data without having already processed it. Bayesian algorithms, like rational people, generally develop more accurate beliefs about a body of data the more they are exposed to it and update their prior beliefs according to new evidence. The issue is what kinds of belief to start with, and how to know that this initial state corresponds to the data without having already processed the data in the manner the algorithm presupposes would be the result.

In Probability Theory, Jaynes quotes Thomas Jefferson. "Ignorance is preferable to error and he is less remote from the truth who believes nothing than he who believes what is wrong." While I would quibble that being wrong and knowing nothing both classify as different kinds of ignorance, this does get to the heart of the matter. In a manner of speaking believing nothing is really to believe in everything, to not assume and therefore to allow events to play out without constricting what one expects by overshooting or undershooting the truth. In essence both the initialization and the bootstrapping problems are rooted in the same circularity. You can't know where to start without actually having an idea about what you're dealing with, but if you knew what you were dealing with then you would already have formed the sought for belief (I.e. Meno's paradox). And even if you come up with something there is no way to know it has any relevance to the problem until you actually apply it, and it would be entirely plausible that one might waste an indefinite amount of time trying out different each one of an infinite number of priors with no correlation until one sticks.

Here things take an intriguing turn. Jaynes does away with the assumption of knowledge priors and instead talks about *ignorance priors*. Ignorance, in a sense, is far from nothing; it stands for the maximum possibilities. To assume a uniform prior distribution of the data, necessarily contains solution within some subsection of it. The sequence of steps involved in the belief updating process

thus start with this "block" and sculpt the solution by narrowing down possibilities as it gains in information.

Jaynes responds to the potential objection that it is impossible to codify ignorance like one can information. "Some object to the very attempt to represent complete ignorance, on the grounds that a state of complete ignorance *does not 'exist'* \[emphasis added]. We would reply that a perfect triangle does not exist either; nevertheless, a surveyor who was ignorant of the properties of perfect triangles would not be competent to do his job. Complete ignorance is, for us, an ideal limiting case of real prior information, in exactly the same sense that a perfect triangle is an ideal limiting case of the real triangles made by surveyors. If we have not learned how to deal with complete ignorance, we are hardly in a position to solve a real problem." (Jaynes 2003, p. 347).

It would seem strange that the ideal prior information would be that which is most ignorant. Common sense would have it that the best prior would be the smartest-the most knowledgeable. For reasons of circularity just mentioned, there is no most knowledgeable prior, no way to obtain a high information state without processing the data. But this is more ignorance in the sense that everything could be possible rather than ignorance as in empty vacuity. If one assumes everything and provides oneself the most leeway, the truth must necessarily be an element in the set of everything. The net is cast wide. If you assume nothing, you have nothing to work with. You haven't cast the net at all. Accordingly, assuming maximal entropy is the "least biased" prior distribution.

There are various technicalities pertaining to this approach to probabilistic inference that are best left to the mathematicians. For example, how to precisely define a value for "complete ignorance"? From an algorithmic point of view? (As it *Probability Theory*.) What is the space and time complexity of a Bayesian artificial intelligence using the MaxEnt initial prior? How much memory would it require and how much time? Given that the entropy of a system is the logarithm of the total number of microstates, the larger the data set the more intractable one would assume taking its maximum entropy as the initial prior distribution would be, as it would require a more resource-intensive search to drill down to the target value. These questions are interesting, but what is of

greater interest for the present purposes is what this says about how ignorance relates to the world.

There is, as I will continue to explain, a profound connection between ignorance and possibility. Possibility, generally understood, is a function of entropy. The physics of it tends to embed and bury this notion of possibility in its theoretical constructs such as a phase or configuration space, obscuring the broader lesson to be learned from it. Nevertheless as Jaynes demonstrated a sort of logical membrane can be stripped from this specific domain of application and generalized. If a form of inference is generalizable, that is another way of saying that it reflects the basic laws of thought. If it reflects thought in the most open-ended form, it must also reflect how ignorance structures thought. Logician David Lewis' quirky but fascinating modal realism, which views possible worlds as actual, can therefore be used as support for the actuality of ignorance. This point will come later however. First we must build out this relationship between thought and entropy.

Given that all changes in the states of matter is an expression of second law of thermodynamics, and given that information theory is modeled on this concept of entropy, and given that the strongest model of cognition is an information (info-entropy) processing account, a clear relationship between these terms emerges. Attempts to connect information processing with entropy are not new, in fact they can be traced to the origin of information theory. What concerns us here however is what all this says about the theory of ignorance.

Thinking is all about change. The physical reality of change is what assumes the need to think in the first place: if nothing changes, why need anything besides autonomic reflex to get along in the world? If there is no change, which is a consequence of thermodynamics, and there is no information, which is a measure how complex it is to describe the displacement of elements in a system, then there is no thinking, which is the process of identifying invariant patterns in the irregular flow of events.

The series of supervenience relations between physical change, information entropy and information processing indicates an ordered relationship between

fundamental physics and the workings of the mind/brain. The mind is at its core an uncertainty reducer because physical entropy is the correlate of informational probability distributions. This parallelism is instructive. It shows that the mind's role is not purely to "represent what's out there" by symbolizing external objects through algebra-like metonyms or intentional objects. It's that the root function of the mind is driven by physical entropy by way of the challenge information disorder presents it. The mind is a direct evolutionary consequence of entropy in the same way that eyes are a response to light, or lungs to oxygen. The mind is the "action site" or the complementary adaptive "groove" shaped specifically to fit entropy. Entropy is the raw substance that the mind works on, and without it, there would be no evolutionary basis for the formation of a mind. Finally the product of the mind is order, it's the reduction of entropy and the sustainment of the structural unity of the organism and its operating environment during the execution of normal healthy behavior. Entropy is the input but through cognitive work cycles, the output is order. Ordered representations that capture the sense of predictability in the toss up of ongoing events.

To get a full grasp on the nature of cognition we must consider all three terms (physical entropy, information entropy, and cognitive entropy reduction) as a system.

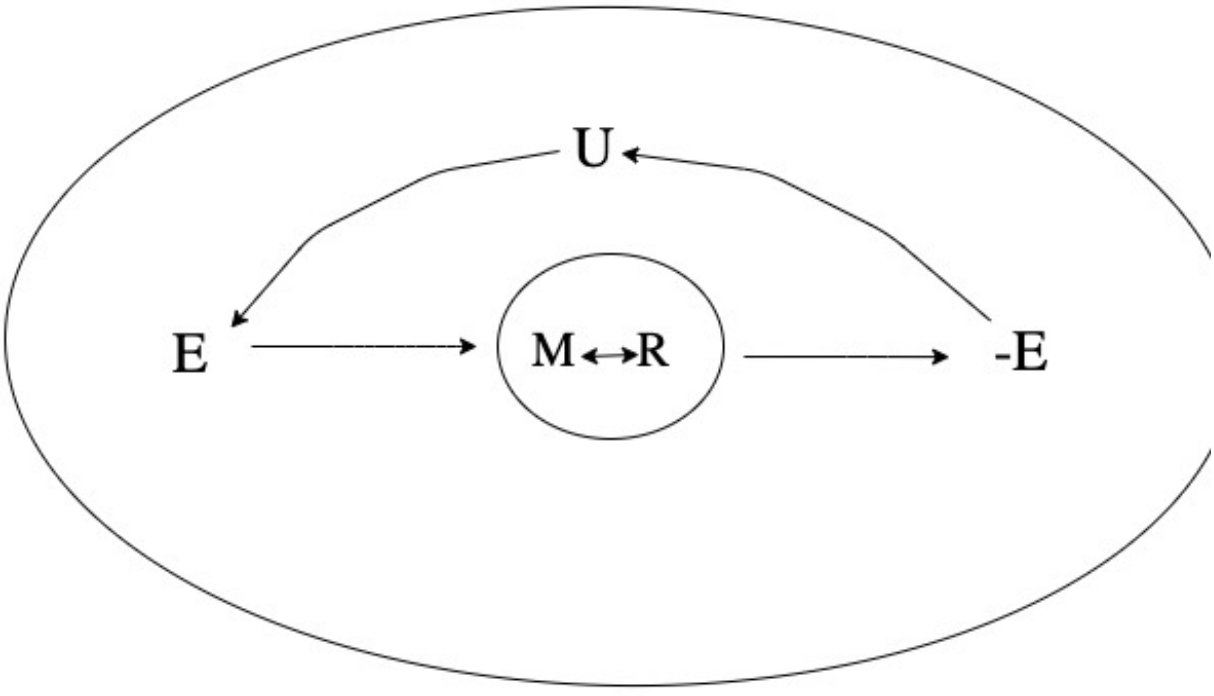

Here, entropy E, serves as input to the mind through the senses. The mind M, entrains representations in response to that entropy. Those representations, R, in turn bidirectionally influence and constrain the mind. The bounded (M + R) forms an enclosed system denoted by the inner circle. The output of mind-representation feedback is negative entropy, negentropy, which reduces the conditions of uncertainty U, further constraining E. If we treat M or (M + R) as the unit of analysis we lose the broader context which constitutes the mind's entire operating environment. If we ignore E, we miss the driver of (M + R)'s cognitive work cycles. If we ignore -E, we miss the product and goal of those cycles.

The general flow goes like this. Information entropy is detected by the senses. M processes this input, and yields R, which identify the stable redundant structures detected within E. This in turn leads to negative entropy, which yields the updated conditions of uncertainty which in turn constrains E. Generally E will become more constrained as this cycle continues and U approximates zero.

The mind is indeed an entropy-wrangler; a system whose primary function is to sort through entropy—or rather, entropy's informational correlate. It achieves its effects by subdividing and excluding categories of entropic noise until stable patterns remain. This is no metaphor. It's precisely the biological rationale for the existence of minds in the first place.

The view that the mind is an information processor is only partly true. Information is not out in the world ready and waiting for sensory pickup. Information is the byproduct derived from raw sensory input played off against established schemas. Information is a residue, boiled down and reduced as a consequence of parsing entropy into invariant representations.

There is more to this analysis of cognitive entropy processing. Because thermodynamic dissipation is the index of time, "time's arrow", all cognition must be situated in time. Uncertainty is deeply entwined with the fact that phenomena are broken up and spread out over time and that each moment makes accessible only a slice of the entire temporal event structure. It is also related to the notion of time asymmetry: if the past were exactly like the future, the past would be a perfect model for the future and as such there would be little

need to think. One would only have to gain enough experience to acquire a representative sample of the future and then apply the past as a measuring stick. As life continually shows, there is always surprises and the future often vastly deviates from expected outcomes. In Newtonian mechanics, the laws of physics apply both backwards and forwards in time, but this does not reflect the reality of entropy and irreversible processes. Once out the cat cannot be put back in the bag, once the beans are spilled and information spreads, one cannot retract it. This phenomenon, arguably, is directly associated to the physics of dissipation. It is quite hard to unlearn what is learned. While memory formation is not strictly speaking irreversible, it is a product of irreversible processes.

The next section will examine the role of time in dictating the conditions uncertainty and ignorance.

All ignorance is conditioned by time. Ignorance about past, including everything up to the last microsecond at which it takes light from nearby surrounding objects to reach you back until the beginning of the universe is one thing. Then there is ignorance of the future, anything that happens by the time the light bouncing off you hits it and onward into oblivion. On a first pass analysis, there is a stark asymmetry between what we can know about the past versus the future. Although the past is done, fully expressed, what survives of it in the present is partial and deteriorated. Even though the present is a consequence of the causal history of the past, even though causally speaking the present is *made out of* the past, the events of the past are represented in the present only in faint impressions, fragments, hints and shards. For a vestige of the past to become known its information state must be preserved against the normal trend of entropic dissipation that tends to scramble and degrade information. Just as the molecular order of physical systems tends to unravel in time, so too does information and with it the knowability of the past. These leftover artifacts are generally the exceptions rather than the rule, the few dinosaur bones that were fossilized in a suitable medium while the rest deteriorated. Good record keeping helps protect us against such information loss, but such methods are imperfect. knowledge of the past is therefore mostly

"filled-in", the product of a pattern-completion exercise in which the whole is fuzzily reconstructed from the parts. The parts here being various causal artifacts of an event having happened, fossils, imprints, displacements, or aftereffects caused by something that was once there.

The future in contrast is a whole different ballgame. Everyone by simple introspection can attest to the drastic uncertainty of the future and that ignorance about it is of a different quality from that of the past. We can uncover details about the past, but the future cannot be inspected. Nothing in our current information base can truly be said to address the future with perfect certainty. The present is causally interwoven with the past and an extension of it. The future instead is broken off from us. From the perspective of the present it doesn't exist yet and is merely hypothetical. Current trends can be projected onto it, but those projections can never enjoy a higher grade of knowledge than mere informed guesses. We can be certain of the relation between past and present in the sense that the past is congruent with it. Nothing in the past happened that would have ruled out the present. While certain changes to the present can modify the future with high reliability (if you break a dish you can bet you won't be eating off it tomorrow), this usually applies to negative facts rather than positive facts. The removal of a concern in the present tells us it will not be active in the future, but the introduction of a new concern usually spells higher uncertainty about the future. The future is fundamentally ambiguous. It is one of the most decisive and important structural determinants of ignorance.

One essential characteristic of thought is that it unfolds both in and over time. Heidegger noted the importance of the experience of time in its ordering of our sense of being and our self-projective relation to the world. We imagine a future we wish to be in, and act to make it real, projecting our current state onto idealized future states and drawing them forth. In this way the only future we can be fairly certain of is our own commitments to it.

The nature of thought in time is severalfold. Not only do discrete thoughts themselves have duration or latency, but they are conditioned by temporal relations to the past, present and future each of which defines certain boundaries of information accessibility and grades of knowability. In other words, thoughts

are about things in the past, present or future, but they are also conditioned by time in that they unfold over time. Time is a medium that shapes the direction and content of thought and which constrains what is knowable and unknowable. To arrive at a deeper understanding of the temporality of cognition and how time constrains knowledge and imposes ignorance it's important to explore all angles of time's impact on the mind.

Entropy is deeply associated with time. Depending on how you look at it, time is either a direct consequence of entropy or intimately correlated with it to such an extent the two phenomena are inextricable. Our sense of time is bound up in the reality of change. Without change, there is no meaningful concept of time. Viewing a clock ticking is viewing a change. Hence the idea that entropy is "time's arrow". The entropy of a physical system always increases and so events unfolding in time cannot be reversed. Everything that happens in time "runs down" towards a single inexorable terminus of maximal thermal equilibrium. In the course of this arrow of time's flying, every potential is actualized and exhausted. The universe's energy dries up as stuff happens.

Any description of time without such a one-way tendency of energy dissipation is incoherent from the perspective of physics. What's of interest here is now the entropic unidirectionality of time and the associated information flow conditions and constrains cognition and imposes certain kinds of ignorance and uncertainty on the mind.

That the subjective perception of time is a modality of cognition is a well-worn Kantian notion. Kantian transcendental idealism views time as a structure of cognition, a condition of possibility for knowledge and perception, a kind of pre-installed mental firmware that organizes all reasoning. We can't picture a thought without presuming time. The subjective sense of time notwithstanding, entropy is a measurable physical quantity and so it makes sense to speak of objective time as well. Factoring in special relativity can make the topic quickly very confusing. Suffice it to say the matter can be clarified by speaking in terms of an inner and outer sense of time.

Concerning inner time, every reasoning involves a temporal- sequential working out of a solution that constitutes a chain of dependencies that rely on

the earlier reasonings. Each step in the chain of reasoning is procedurally reliant on what steps before it in the sequence and therefore in time.The same can be said about the use of memory, prior states of the system influence its future state. Reasoning presumes time.

Reasoning arguably has this sequential form because the nature of time determines it. The serial nature of reasoning unfolds within sequential outer time and repeats it as inner time. The world presents only so much data at any moment. And the mind has limited storage and processing capacities, meaning that it has to load-balance and distribute its reasoning over a protracted sequence rather than engage in a parallelistic working out of all the problem parts simultaneously. Time slices up what is knowable into different apportionments depending on where thinking stands in relation to each moment of its unfolding.

All or nearly all of the future and much of the past are uncertain or unknown and perhaps unknowable. We can have *beliefs* about the future, even true beliefs, but we cannot know we know until the predicted event happens. This is perhaps one of the defining facts of life and the core problem of decision making. Time carves out the contours of ignorance in the most general terms. There is ignorance about the past, the present, and the future, and distinct conditions that apply to them.

From the standpoint of the availability of evidence there is always a temporal incompleteness relative to past and future. From the perspective of the present, the past informs us only through fragments and traces, its full information state having since dissipated. The future presents nearly an impenetrable wall of ignorance as no information at all is communicated by it. If we could peer through the dimension of time like we do space much uncertainty would be negated. We would perceive a salient, unified Einsteinian “block spacetime” where everything, past, present and future stands frozen in place. To find out what would happen in the future, we’d merely have to direct our attention toward the region of spacetime in which it belongs. We would experience a “universal present” where the future and past is just as *here* as the now. Instead time perception is truncated into what amounts subjectively to a “spotlight” of the present dragging across a series of infinitesimal moments going from past to

future. Within that spotlight we can detect relics of the past, and we can reason about the relation of those past relics to the present and what that might anticipate about the future. The future in contrast, seems totally cut off, it cannot communicate to us in any way. So we are limited to imperfect reconstructions of the past, evaluation of the present, and incomplete probabilistic inferences about the future based on the surmises gathered from the former two sources.

From the fact that we must factor together two incomplete information sources, past and present, to reason about the future goes a long way to explaining why the future is inscrutable and can only be anticipated in broad, tentative outlines if at all. An unknown amount of information about the past has been lost, and only so much can be found out in the present before it immediately becomes the past. Together, it's unclear how much what we can surmise in of past and present tells us about the future. There simply isn't enough information to perfectly predict the future in exact detail, and this is an inevitable consequence of the time asymmetry predicted by entropy, the fact that the future is more disordered than the past. Or if there is enough information, it is not feasible to collect it all. Forecasting is such a tricky practice because of these constraints. This outer time, the objective time that defines the temporal structure of limited information flow from past to present, cuts off abruptly in the current moment leaving the future a gaping, speculative abyss.

The present is a very slight pinpoint. Rarely can all the relevant information be brought to be bear all at once. There is only so much thinking one can do in a given moment or sequence of moments. Thinking is therefore temporally extended in the sense that thought takes on more content over time as new information flows in. All ignorance is twicefold. We may not know how to solve the problem, and we may not know how much more must be known to determine when sufficient information has been acquired to know we can solve it. This first order and second order ignorance unfolds in time. In other words ignorance is bounded both by the individual requirements of the problem and by the lack of criteria by which to measure the adequacy of our knowledge. It's one thing to not know how to solve a problem, it's another thing to not know what needs to be known to solve it when you can't. In time new pieces of the puzzle

may fall into place and second order meta-ignorance will be corrected, which allows for first order ignorance to be in turn rectified.

Because entropy is the expression of time it is clear why time is essential to the analysis of the mind given that the mind is fundamentally an entropy-reducing utility. Seeing that entropy and time are two sides of the same coin, this would seem to suggest that cognition is time-reducing too. That's a strange notion, but in a way it makes sense once we locate the essence of time in change.

The result of cognition is to form invariant representations. To identify and *freeze* perceptual regularities in neural codes. In this way representations are a means of detecting what doesn't change about recurring patterns of change. knowledge is a move toward stasis. Without change there is no time. So to establish sound representational knowledge is to arrive at a sort of timelessness. A feature of the world is then perceived as stable when before it was indistinguishable from the chaos. At which point we don't have to think about that now static part of the world. It becomes second nature to unconsciously depend on and navigate by its lights. So the result of (true, accurate) thinking is to counter time, to detect sufficient regularities in the flow of change that a whole sequence of changes is itself seen to no longer change. Indeed this connection between time and thought may partly explain why our perception of time changes as we age or why time flies or slows to a crawl when we're having fun or at the dentist. When processing is fluid and frictionless, awareness of time effectively disappears, and it goes by fast. In contrast when there is thinking to be done, time slows and we must go through the paces consciously and deliberately.

The goal of thinking then is to stop thinking, to arrive at a conception of invariance. The very fuel of thought therefore is *variance* which is another word for uncertainty/entropy.

Now having described the concrete relation between cognition and entropy let's explore some other physical phenomena in search for clues to the nature of ignorance, knowledge, and belief.

In the analysis of the connection between information processing and thermal entropy three concepts are traditionally employed in the literature: Maxwell's

demon, the Szilard engine, and Landauer's principle. The former two are essentially analogs of each other, so it's best to explain them together. Landauer's principle leans more to toward computational physics.

Maxwell's demon is a classic thought experiment proposed to speculate about how a hypothetical exception to the second law of thermodynamics might be implemented. For our purposes, it's only necessary to give a brief sketch of the idea. Picture a rectangle that contains a a gas at uniform temperature. That is to say, fast moving and slow moving ("hotter" and "colder" particles) are present in equal amounts throughout. The box is divided in the middle by a partition with a gate that when opened allows a particle to pass from one side to the other. At equilibrium, both sides of the box are essentially identical at the macroscopic level and equally randomized at the microscopic. Presiding over the gate is the demon, who possesses extraordinarily acute senses that enable it to see when a faster or slower particle is on course with it. The demon's job is to watch for fast moving particles, which would necessarily be tailed by slower particles, and rapidly open and shut the gate so that only the fast particle passes. On the warmer side, the demon does the inverse, blocking the passage of a slower particle every time it intercepts the gate. Repeating this operation indefinitely, the demon would ultimately divide the rectangle into into a hotter and a colder side thus contradicting the second law.

The implication here is that Maxwell's demon its something like an information processor. It effectively asks a yes or no question every time a "bit" approaches its gate, structuring a low information system into a high information system. It engages in a *negentropic* structuring process, ordering the environment by taking it from an unpredictable to a more predictable state. The information connection was made far more explicit in physicist Leo Szilard's engine concept.

A Szilard engine operates a bit differently from Maxwell's thought experiment. In principle it consists of a box with a sole particle, which is exposed to a heat bath which causes it to fluctuate in position randomly. The partition which separates the box into two parts is not stationary like in Maxwell's example. Rather it is capable of moving to the left or right and tracking the

location of the particle. The demon in Szilard's conception is detects the location of the particle, and moves the partition closer to it. Also important to the engine is the notion of pressure and a piston. In simplified terms, when the particle strikes the partition at this point, the pressure which is built up in the remaining area between by the partition and containing the particle can be used to exert useful work by being transferred to the piston. The portion then moves in the opposite direction of the particle, releasing energy which can be channeled into the actuation of a mechanism. When it reaches one edge of the box, it reappears on the other side, allowing the process to repeat indefinitely. There are a few more technical details involved, but in broad outline this captures the idea. By detecting where the particle is located the demon can move the partition towards it, making free energy available that can be extracted to do useful work.

Thus by tracking a single bit of information the engine can extract free energy by "chasing" the particle. This tracking does come at some cost of entropy, but the fundamental point is that information processing and thermodynamics can be linked concretely. By measuring the location of the particle, the demon is able extract work. There is a problem however. As Wolpert and Macready famously (1995) put it, there is no such thing as a free lunch, every computational operation always comes at a cost ([https://www.semanticscholar.org/paper/No-Free-Lunch-Theorems-for-Search-Wolpert-Macready/93a7f5b7f51622430a4a4d4002152b277e4be470?p2df]). In this case, the cost is in heat produced by the requirement for the demon to rewrite the information about the particle's location each time the cycle repeats. Thus the Szilard engine is no perpetual motion machine. It expels heat energy whenever it performs an irreversible logical operation such as rewriting its memory store. It could never be self sufficient and self powering. This finding was established by Rolf Landauer in his landmark 1961 paper *Irreversibility and Heat Generation in the Computing Process*. ([http://worrydream.com/refs/Landauer%20-%20Irreversibility%20and%20Heat%20Generation%20in%20the%20Computing%20Process.pdf])

At this point the theory of computation and physics blend. While this chapter of the book attempts to analyze them separately, the fact is they are inextricably

bound. Nevertheless, these findings bear important implications for the theory of ignorance and dividing them into arbitrary partitions does not serve to bring those insights to the surface.

Landauer established that logically irreversible operations, such as erasing a bit from memory, must correspond to a proportional loss of information either within the device or its environment. This point can informally summed up as "take one, leave one." A local reduction in entropy entails a global increase (*k*T ln2 of an amount to be precise). The uncertainty reduction achieved by information processing does not actually eliminate it from the world, but shunts it off into the environment or into the system itself supposing it did not have an appropriate heat sink, leading eventually the fragmentation and disorder of its memory storage. Landauer's findings also achieved the important result of firmly associating information and physicals states mathematically (Wright, 2016 [https://www.mitpressjournals.org/doi/pdf/10.1162/POSC_a_00194]), which helps shore up claims that ignorance-as-uncertainty has an objective basis.

It's natural to ask, what does does this all have to do with epistemology? The goal here is not to give a textbook overview of these scientific principles. There are far better resources on the topic. Rather, the goal here is to take these findings out of the well-defined theoretical cubbyhole they were developed in and expand them to encompass broader horizons. There is indeed a deep connection between these concepts and the wider concerns of knowledge and ignorance. Admittedly, the conclusions to be drawn are conjectural, but that is only because they are novel. From what I understand there has been no prior work to assimilate these findings into epistemology.

In no trivial way are human thinkers much different from the particle obsessed demons mentioned above. The key difference is that it is our own minds which constitute the box. Our critical faculty is analogous to the demon, and the information state the disordered helter-skelter of the undifferentiated initial distribution. In a very real sense cognition is preoccupied with separating the proverbial wheat from the chaff, separating truths from falsehoods and building up a concentrated representation of what is true on one side and what is false on the other. Through a self-organizing mural exclusion, those propositions

which are deemed to be true and those that are deemed to be false although initially entangled and indistinguishable, are tied together by logical associations and pulled apart, the truths flowing to one direction and the falsehoods dragged to the opposite ends. In the process of clarifying our doubts, we play the role of such a demon, evaluating each proposition-as-particle and saying, "do I have a reason to think this is true, if so, permit it entry to the truth side. If not, keep it out."

How does this relate to Landauer's principle? The brain does indeed consume significant amounts of energy in the course of its operations, constituting only 2% of body mass but burning up an estimated 20% of its energy ([https://www.scientificamerican.com/article/why-does-the-brain-need-s/]). And it runs hot. But epistemologically speaking we are not concerned with the production waste heat. It's mainly irrelevant, although ignorant people do put out quite a bit of hot air. What I am about to suggest is not a direct entailment of the science from which it draws inspiration, but which has its own logical consistency and reasons for supporting it. It is somewhat related to the so-called law of intended consequences, or to put it cheekily, just as with more money more problems, more knowledge more ignorance.

At first this makes no sense. More knowledge, less ignorance, simple as that. knowledge, and the things it allows us to create, opens doors. It creates more possibilities. But according to our dictum, an increase in possibility corresponds to a proportional increase in information entropy—more ignorance. The knowledge used to construct the automobile created much of the landscape of modern civilization. It also increased new demands and requirements, and introduced new problems such as pollution and dependency on fossil fuels, which entailed more ignorance dependent on that knowledge needing to be resolved in order to maintain its feasibility as a technology. knowledge of nuclear fission directly entailed the mishaps that led to the Chernobyl disaster. We have a tendency to think that there is a linear proportional relationship between the increase of knowledge and the decrease of ignorance, but this is not so. Certain knowledge "activates" regions of ignorance which were once irrelevant. A medieval farmer was ignorant of nuclear physics, but they were also mercifully

ignorant of the prospect of nuclear war. Thus the more you know, the more you realize how little you know. This is partially because you have gained a better perspective on the epistemic landscape, but it its also because *knowledge can activate more ignorance*. knowledge can potentiate ignorance, forcing what was once remote possibilities to the fore.

Today we find ourselves in the baffling situation where drastic exponential improvements of information technology, tools intended to help us process data and make the world more predictable, have made the world instead much more complex and difficult to understand. That's because those instruments generate their own data, adding that to the pile that was already there. And they brought with them their own complexities which have to be factored on top of the preexisting complexities. And as these systems grew and scaled, more technologies needed to be developed in order to handle and cope with the the problems they generated, and so on. To the extent that knowledge engenders an increase in complexity it also does an increase in the potential for being ignorant in new ways that count. While knowledge often allows us to simplify, local simplification may lead to global complexification. The invention of cars certainly simplified the task of getting around long distances, but as it turns out this problem does not exist in a sealed bubble but interrelates with a web of other latent factors which were only brought into play once the initial transportation problem was solved. Congested highways, car accidents, you name it. The reason for this is that although we are less ignorant once the discovery has been made, we are nevertheless now *ignorant of the world which has been newly changed by that discovery. *

One might counter this by saying that the total ignorance remains constant. The dangerous side-effects that nuclear engineering introduced to the world were always there in an abstract sense. In the fullness of logical space there is nothing new under the sun. Regardless of our awareness of it, there was always some set of possible worlds where those consequences held true and our timeline just so happened to intersect with it. In response, I would say that not all ignorance matters the same. A world without the discovery of nuclear physics is a world where ignorance of its implications not only doesn't matter, but is in

some sense undefined, because only a world in which it has been discovered do those implications carry significance. I do not mean to suggest that knowledge, specifically of an engineering variety, invariably causes more problems than it solves, or even that there is a constant correlation between the increase of knowledge and newly relevant ignorance. I do mean to suggest that thinking and its consequences carries with it a burden of irreversible cognitive "waste heat." By changing the world in a manner corresponding to an idea, the resulting disruptions shake up events in such a way that sometimes gives us more of importance to be in the dark about.

Even in the case where those close to the discovery see its possible implications, for example the Einstein-Szilard letter to President Roosevelt informing him of the possibility of creating nuclear bombs using the newly discovered technique of nuclear fission—in accordance with the principle that *a valid generalization does not give you insight into the application of its rule*, this realization does not tell us how that general truth will eventually find expression. Powerful knowledge results in unintended consequences, especially if that knowledge has practical applications. But even in purely theoretical realms, new findings restructure the space of ignorance and affords us new riddles to deal with in order to reconcile that incomplete knowledge with what it entails beyond its scope.

This study of heat physics equips us with several tools and key insights which illuminate the problems of ignorance. To take it a step further I wish to broach the concept of knowledge gradients, a difference between lower and higher information across a threshold which anticipates the possibility of learning on the lower information side of the divide.

A recurring theme is the idea that there are profound parallels between thermodynamics, information theory, and epistemology and the psychology of learning. Rather than a merely superficial analogy, the argument I propose is that there are deep supervenience relations connecting these phenomena. Learning — on the most fundamental computational level—is information entropy (uncertainty) reduction, and information entropy is the correlate of the

disordering effects of thermal dissipation. The scientific literature relating thermal and information entropy is rich and dense, but less has been written about how that relationship might feed into epistemology. Understandably, as both thermodynamics and information theory are deeply quantitative fields and Shannon's equations share clear mathematical parallelisms with thermodynamics while epistemology lacks this quantitative basis. Indeed, placing the Boltzmann's and Shannon's entropy formulas next to each other shows the unmistakable similarities of their logical form. Both measure the probability of a given element (microstates versus messages) within an ensemble (a gas versus a message space). Epistemology however has remained rather parochial in terms of assimilating these findings into its conceptual toolkit, much to its loss.

The key challenge is to connect these deep physical and informational concepts to an understanding of knowledge and ignorance. One thing is clear from the outset: although there are supervenience relations defining ordered relationships between these three categories, the principles and operations that apply to each are not the same. For example gain in information necessarily means a gain in thermodynamic heat, because computation requires physical work (Barato, D. Hartich and U. Seifert, 2013). So stating that learning decreases thermal entropy is false even though it correlates with a decrease in information entropy locally. Nor, of course, do we simply derive knowledge from reading the temperature. Knowers are not thermometers. Rather, thermal dissipation implies increased randomness in the distribution of environmental states which correlates with increased information complexity. The information measures the probability distribution of that outlay. On a fundamental computational level the information processing that underlies sensory perception functions to equalize the mutual information between the environment and the agent (Stamps, 2004). In other words, the function of perception is to bring the internal information of the sensory system into equilibrium with external information. When an agent perceptually explores an environment, they generally continue to do so until the environment is fully "mapped out" or some goal state has been achieved. To learn is to approach homeostasis with external information.

The reader may notice that my usage of "equilibrium" differs from that used in physics. Equilibrium in physics typically denotes a maximally disordered, high entropy state. Yet this difference is explained by the obvious reason that when someone learns something, it does not destroy or lessen the information contained in the information source. Again, it is misleading to assume that the parallelisms across these domains are undeviating and foolish to assume what applies to particles in a gas must also apply to human minds communicating and learning from each other. A teacher does not "expend" their knowledge by teaching it to students, if anything they reinforce that knowledge in themselves by repetition and reiteration. The increase of mutual information between the internal and external environments is therefore not inconsistent. It is the mutual information that tends to evenness.

The environment is an information sink as much as a heat sink, to which the prevailing attitude of the organism is one of information deficit. "I need to know more, and the only way to find out is by exploring the world" is the governing principle. Because there is an asymmetric disparity between internal and external information, learning is required to minimize that discrepancy. The same relationships hold between an agent and knowable unknowns. This is where the concept of knowledge gradients comes into play.

Consider the definitive example of a university lecture. The premise behind the whole affair is that the lecturer knows what the audience does not, or at least is reminding them of what is not fresh on their minds, which nearly comes to the same thing. Otherwise nobody would sign up for the course. Such a setup specifies a directional knowledge flow from the more knowledgeable to the less knowledgeable (more ignorant). A gradient—information asymmetry—defines the relationship between teacher and student. This need not necessarily be a fixed one-directional relationship, students can and do provide insights to the class discussion that the professor did not anticipate, but in general the overwhelming tendency of information flow is from teacher to student.

When two or more communicants all know the same information, there is less reason for them to converse. "Tell me something I don't know!" one might say. The parity in knowledge makes redundancies more frequent and

communication more wasteful. (Although the may converse in order to establish the mutual knowledge that everyone is "on the same page" and resolve the uncertainty that they do in fact know the same things. Even if they know the same things they do not know they know the same things until they have aired their knowledge to each other.) While various parts of the shared, collectively distributed social information system (common knowledge) are called up given certain relevant contextual cues, as when group attention must be directed to one focus at a certain stage in a plan, generally we communicate with each other when we need to find out something new or tell someone what they need to know. We socially plan and converse until the cooperative group is "on the same page." In a cooperative game, such coordinations become balancing acts, in which an over-concentration (lack of diffusion) of information in one subset or member of the group negatively impacts another, implying feedback loops and dynamic equilibria that constantly readjust between participants so that people are exposed to information on a need to know basis at the right time and place.

In adversarial, zero-sum games, the goal is not to equalize group information but to acquire an information advantage. The object is to effortfully restrict the spillage of classified information and to gain power over the more oblivious party. Governments establish tiered security clearances and classify documents so that knowledge of them is not leaked to the wrong people. Corporate firms covetously guard trade secrets and intellectual property to keep an edge over competitors. A liar withholds the truth so that the person lied to believes and acts in ways that are consistent with the lie. It is no coincidence that we speak of "leaks" of sensitive government information, the diffusion of information is more than just tangentially related to a kind of fluid mechanics of dispersion and spread. "Information wants to be free" said writer Stewart Brand, which is a way of saying that there is always a probability that we can't "keep a lid on it" and that there is always a risk that information spillover will propagate through social networks unrestrained.

In either the cooperative or adversarial case, there is a discrepancy in gradients that is the structuring principle of the entire dynamic.

When nobody knows the answer, attention turns to the general environment

and investigation shifts to the world at large. We stop being each other's gradients and instead we turn to the world for instruction. As rule, whether it is between each other or between us and the world, a disparity between relative knowledge and outlying ignorance drives learning. Agents must continuously couple their interactions with changing environmental conditions. Because of this need for coupling and information updates, knowledge cannot exist without explorable ignorance, a discovery space of available information.

In computer science a widely used method called gradient descent is an optimization strategy for artificial neural networks. Essentially, the gradient specifies a slope that expresses the values of a differentiable function—a function that measures the derivative or rate of change of an input. The higher the gradient, the steeper the slope, the faster the model learns. If the slope is zero, that's the same as saying there is nothing to learn, and learning stops. The purpose of a gradient decent algorithm is to minimize a cost function by finding a local minimum solution using the fewest steps (Ruder, 2017). While this context of use is not precisely the same as what I describe, they are similar enough to be instructive.

The world at large is a very, very steep slope. From birth, learning happens rapidly and automatically because there is everything to learn. It is easier to be original in the past rather than the future. Eventually this slope levels off and plateaus and the learning rate slows, although objectively it never reaches zero. The true picture of learning is of course more complicated and nuanced than the idealization presented here, but the fundamental governing principle revolves around gradients. Knowledge gradients, and the parallel ignorance it works against, is a useful concept for understanding the intrinsic relationship between them.

They key to this argument is that there is an objective basis to beliefs, their structure, and their level of correctness, despite shifting appearances of unassailable subjectivity and opinionation and the problem of other minds and their beliefs.

Unpredictability and uncertainty arise in science in many manifestations,

most notably perhaps in quantum mechanics. According to the uncertainty principle the mass and velocity of a particle cannot be precisely determined in tandem. The determination of one or the other is an exclusive OR proposition.

Enveloping the particle prior to measurement is a probability shell of its possible locations, which cohere upon observation into definitive fact about position and momentum.

In a similar way one can picture a belief state as analogous to the particle. A person deliberates on their beliefs and consolidates it through a process of sedimentation, in Marleau-Ponty's parlance. Prior to arriving at this determination and subsequent commitment to the belief, the subject is working in what might be described as a probability cloud or haze of uncertainty of their own. The world— and their context in it—ought to inform the belief if the target of the belief is reality and not simply relatively useless fantasy. Beliefs, as the pragmatists again argue, are operative—they stretch out and connect with the world via the actions they entail, and via communication and deliberation with it and other people, they connect to other beliefs. The act of acquiring a belief collapses the probability or information-theoretic free energy into a determinate structure.

This analogy about the mapping the case of belief in the context of uncertainty to the case of measuring quantum particles, even if it is simply an intuition pump with no objective basis for its verisimilitude, helps capture the problem of how to ground phenomena of opinionated discourse.

The comparison of a mind to a particle is almost laughably reductive but they key is in the difference that makes a difference: a particle has no memory. The fact that minds sedimentize— acquire persistent and cumulative structure over time in a productive and instructive manner is what stands it apart. This organizing principle of representational structure acquisition feeds into the predictive coding paradigm in which the agent iteratively develops a greater capacity for "honing in on" an accurate predictive appraisal of its working environment through sensorimotor testing, familiarization and adjustment (Clark, 2016).

Yet in acquiring that sedimented structure it's very much a garbage in,

garbage out, goodies in, goodies out scenario. The recognition of good and bad information sources is a pedagogic standard. And through reflection it's clear that not all thoughts an individual has are necessarily beneficial, and yet those thoughts are still *legal* in a strictly logical rule based sense. People can be elaborately mistaken, even in some sense productively mistaken. For a computer, malformed code will fail to compile. But bad thoughts are not bound to the same strict rule. At best, they will be confusing and misunderstood. But what is going on here in these cases?

Consider the connection between determinism and certainty. That we can predict that a process will continue to execute in a selfsame manner indefinitely until acted on by some external agency, such as Newton's first law of motion, is the paradigm of certainty. As Prigogine (1997) articulated, the existence of irreversible process places doubt on the deterministic picture of physical reality. The world may be inherently indeterministic, which means that knowledge is fragile. knowledge that an event can only happen in one way by strict necessity is the only circumstance where maximal certainty is valid. Whether this world truly obeys such necessity in broad strokes is unknowable, because we can never reach the status of Laplace's Demon and acquire infinitely detailed information into the world's state. In his disturbing and profound novel *Blood Meridian*, novelist Cormac McCarthy wrote:

> The universe is no narrow thing and the order within it is not constrained by any latitude in its conception to repeat what exists in one part in any other part. Even in this world more things exist without our knowledge than with it and the order in creation which you see is that which you have put there, like a string in a maze, so that you shall not lose your way. For existence has its own order and that no man's mind can compass, that mind itself being but a fact among others.

This passage grasp an important truth. As mere fragments of the entire universe we are destined only to ever experience and know a fragment of all eventualities, never to adopt a God's eye view, *sub specie aeternitatus*. We might theoretically understand that the "world is all that is the case" as Wittgenstein put it, but we will never receive news as to what "all that is the

case" is. That the world is all that is the case, despite hinting that all facts frozen in a kind of matrix of certainty, flies in the face of subjective ambiguity as to the layout and distribution of those facts. Our own experiences is like a mere tangent line grazing the outer surface of a deeper, unexperienced totality of which in our decades of life we can only hope to sample an infinitesimal fraction. Most of those facts in the manifold of logical space are occluded from us. Does this mean that the whole "super-fact" of the world is one crystalline, absolute essence that is only unknown because it is not in our acquaintance? That in other words, there are no ifs? Is everything given all at once, even though we are blocked from seeing through to it and have to wade through thickets of subjective uncertainty? If so, then uncertainty is merely a mental state, if not, then what we respond to when we experience to uncertainty is a mental representation protean fabric of events in themselves, or we are ignorant of things because they do not exist yet and so their status as facts is undetermined.

It's natural to conceive of the world as a set of potentially branching, counterfactual alternatives, and to pick and choose which to believe and commit to of those we wish to see happen or to avoid. The universe is one great "could be." We say to ourselves, "if that happened, then I would have done that, but it never happened, so I did this." To want something is on some level to picture the missing object of desire and then to act in such a way as to bring into actuality. We intuitively think of the world's event structure as flexible; that within its events are latent potentials that can be actualized. Fatalism is stultifying and depressing, it preempts action and therefore makes itself self-fulfilling. Our actions bring facts about which otherwise would not have come to pass. Things can turn out differently given different conditions. Indeed this is baked into the grammar of modal language "could, might, would, can." In fact, it would be foolish not to think in terms of contingency and throw caution and planning to the wind and act as if only one outcome was possible in any event. Insurance policies, sophisticated financial instruments and investment strategies, and general life planning and preparations all follow this principle that there is a multiplicity to how things can shake out. It's a normal feature of decision making to parse a course of events into its moving parts, and map out the degrees of

freedom of each part and how if they swayed one way or another that could change the whole flow of events.

The distinction between probabilistic and strictly deterministic causality has drastic consequences for how to establish a proper foundation for epistemology. Certainly this has been widely recognized, which explains the modern emphasis on modal logic in much cutting edge research. It might be argued that there is no such thing as probabilistic causality or possibility, that causality implies necessity. But it may also be the case that things are necessitated to be uncertain. Given that we can't help but wonder if one or the other of these are the case, modal reasoning forces itself on us inescapably.

It is an essential ontological commitment if it can be said that modal language is denotative, it describes the way the world works rather than how our minds are forced to conceive of it. The question of the ontic status of objective versus subjective possibility is critical because it decides whether the world is inherently uncertain as a consequence of its capricious nature or if uncertainty is solely a property of undecided mental states. While admittedly there is no clear way to prove either position, I argue the former, with some provisos, hence why I seek to ground epistemology in a framework of uncertainty rather than false certainty. The main reason for my commitment to objective uncertainty has been explained elsewhere. If we can't help but be uncertain, then there is no reason to treat that uncertainty as if it weren't an objective fact. After all it appears everywhere we look, and we cannot banish it no matter how we try.

As usual calling on physics is a good heuristic to look for some insight into the issue of the plausible objective reality of indeterminacy. Within physics a formula called the path integral formulation generalizes the principle that a particle's trajectory can take an infinite number of possible routes. The formulation describes a probability density or *sum-over-histories* of all the possible future states of the particle subsequent to its state at the time of measurement. The particle is said to exist in a state of superposition in which several probable trajectories hang in suspension. If the particle can be in many possible states, the idea is to consider its most average state to be the sum of all states. The path integral is thought of as an objective way of reasoning about

quantum states. It provides a way to find the upper bound on the indeterminacy of the particle but "drawing a circle" around all its potential state configurations. The path integral therefore gives us as close an empirical basis for the objective existence of possibility that we are likely to find.

"You never step in the same river twice," goes Heraclitus' old chestnut. If no state of the world ever exactly repeats, then every if-then conditional trigger only specifies a strong likelihood that the following outcome is strictly determined by it. How can we be anything other than ignorant? We may create near identical conditions and see the same result, but this does not always mean that the same result will appear given those conditions (hence the problem of induction). Quantum physics tends to describe its objects as if there is a definite physical field of potentialities surrounding them. The quantum object is enshrouded in ambient particle states that suddenly decohere into a single measurable result. It's natural to speculate that this same principle might apply to macro-scale classical objects though a so-called universal wave function.

To understand the universal wave function and how its relevant to the problems at hand its necessary to understand the conventional, widely acknowledged Schrödinger wave equation. The universal wave function is an audacious generalization of the quantum wave function to the entire universe, blowing up its inferences about the measurement of subatomic particles to ultimate scale.

To understand the Schrödinger wave equation, it's prerequisite to take into account certain main concepts. A subatomic particle is described by several eigenstates, which specify the probability distribution of the possible outcomes of a particle's positions. The total set of eigenstates states the particle is in is circumscribed by its initial energy level. Its energy in other words limits where it can go within certain bounds. The wave function describes where a subatomic particle *can be*. It gives the distribution of its likely positions. A quantum state is said to only be definite once a measurement, all called an observation, is made. According to quantum theory, the particle exists in all its probable positions at once, or rather, in an undecided superposition. The particle ultimately comes into contact with the external world, which *collapses* its eigenstates from many to

one; this is what is meant by observation or measurement. The wave equation thus lets us bracket the total state of the particle and to describe its transformation from a probable to definite state.

The Schrödinger wave equation is philosophically interesting for a number of reasons. Possibly the most interesting reason of all is that it asserts that probability has an ontological reality and is not merely an expression of the degree of subjective doubt about an unknown fact. It shows however that the upper bounds of what's possible for a system is absolute. It's bounded by the energy level of the particle. That would seem to mean that we can be certain of a limit to the uncertainty about the particle's evolution. Such a formulation gives a picture of the knowable constraining the unknowable. The fuzzy indeterminate state of the particle system still evolves according to a predictable trajectory despite its inner content only being specifiable as a probability distribution.

Many of the pioneering early quantum physicists were disturbed by these indeterministic results, with Einstein quipping that "God does not play dice." If the world is indeed probabilistic in its fundamentals, in the sense that the particles that hold it together embody probability as an ontological actuality, this would carry profound implications for the relationship between uncertainty and the mind. It gets weird here. As it would seem to suggest that uncertainty is at the root of reality and thus the mind directly touches it when it processes uncertainty.

It's about to get even weirder when we take into account the universal wave function. The universal wave function was proposed by physicist Hugh Everett as the wave function of the entire universe. If you take what was just said about a quantum particle and map it to the totality of what physically exists, then what is true about probability in the first case also applies to the second. The whole universe then can be said to exist in a superposition of all its possible states, only for one to become actualized in the version of the world we observe. Except on Everett's interpretation, when the universal state enters into measurable decoherence it splits into separate universes which all exist equally. A coherent quantum system is one which has overlapping wave phases. So long as it doesn't interact with the external world, a single system's coherence can be maintained

indefinitely but is unobservable. At the moment of decoherence, when one actuality manifests from the phase space of probabilities, Everett argued that indeed there was no collapse of the wave function but that each possibility branched out into a parallel universe. Instead of a set of fuzzy indeterminate states consolidating into one actual state with the rest discarded as merely might-have-beens, the object takes all trajectories simultaneously. The actual state we see is just the one that played out in our universe and meanwhile all the other eventualities took to their own. The resulting Many Worlds Interpretation of quantum mechanics provides interesting inroads to Lewis' modal realism from a physics perspective. But what is of principal interest here is that it suggests that the same rules of probability and possibility governing the quantum world might apply to the classical world of everyday, macro-scale observable objects.

There is a curious parallel between how physicists describe the quantum world and how we think about our own position relative to the unknown. When we are uncertain about a proposition, our own cognitive state is in much like a moment of superposition. Without some act of measurement or evidence—an influence from the outside world—we might consider anything within the realm of possibility plausible. And typically, our gut is to go with the average, the sum over histories. Only by seeing the fact come together under our watch in tandem with the evidence does more certainty about an outcome congeal before our eyes. We can't escape thinking in terms of what-ifs, and this is a direct consequence of the pervasive universality of ignorance. The quantum probability shell is a model for our own thinking about the phase space of events, that is, the combinatorial aspects of the world which tells us that things can play out in numerous ways granted different arrangements of factors. Again, that alternatives are conceivable is the definitive of the concept of uncertainty. Without it, there is no choice but to be certain: either certain that we know the outcome or certain we do not.

The true test of thinking however arises when a unique event awaits us, one for which there no backdrop of experience to assist us in simulating its superposition of eventualities. Here enters the problem of induction, the most interesting of problems. Anomalies therefore present a special case for studying

the logic of reasoning under uncertainty. Later we will tackle the problem of induction head on, but its relevance to all reasoning under uncertainty is pervasive.

When thinking about the possibilities of an event is this modeling strategy aligned with how the world works, as one might expect of a purely *representational* engine—a brain that works by accurately reproducing the causal structure of the world in order to successfully anticipate and preempt it for its goals? Or is it a merely constructed, conjectural overlay related only to the brain's information entropy as it works out the one and only necessary answer to the one and only necessary outcome? Is this form of cognition more like sense perception, the detection of what's there, or is it more of a flurry of guesswork, populated by a set intentionally inexistent imaginary objects in which we try to find the needle of fact in the haystack of guesses? If Lewis' admittedly rather out-there modal realist take is true and possible worlds exist, then representing events in terms of possibilities is indeed a type of sense perception. Everything we come up with is a fact in some possible world, and thus, the task is to identify which fact inheres to our version of events out of the multiverse of all facts ranging over all worlds. Otherwise, the brain's processing of branching alternatives is a sort of entropic exhaust heat, a byproduct of the procedural derivation of a determinate, actual outcome which must first slough off successively approximating variants of the true proposition until the final answer is obtained.

In the common sense view of perception it is held that our sense *track* an externally adjacent stream of phenomena. Senses attach to the events of the external world, and if we perceive a sensation accurately it means something in the world is there or has changed. There is no sense organ for entertaining possibilities other than the neocortex as a whole. Possibility is not regularly interpreted as a sense modality, and because we cannot detect other worlds, there is no grounds for deciding with certainty that it is. On the other hand if the consideration of alternatives is rather a kind of cognitive waste heat, a necessary byproduct of doing the dirty work of refining a solution, the conception of possibilities is the working through of uncertainty reduction as we arrive at the

boiled away remainder comprising a final decision to commit to a proposition.

In everyday experience anyone who has ever solved a problem can attest to viewing a scheme of events and seeing some way for it to change that resolves the issue. This change was not immediately "broadcast" by the event scheme itself, but is a novel insight projected onto it. This ability to insert a "missing piece" into a scene which otherwise does not indicate it is the key to creative problem solving. At some point our ancestors must have innovated a capacity to see that things could be otherwise than they are, thereby providing the fundamental principle which has guided all human progress. To see something new in the world that isn't there, but which is still consistent with the rules governing objects and their relations, and then to act to realize it, is the bright mark of practical intelligence.

This discussion of the provenance of counterfactual possibilities wanders too far afield for the present purposes. Later a more detailed analysis of ambiguity will be ventured. As it stands, two and only two options can be sharply delineated: either the cognition of alternatives is a product of the direct connection between representational cognition and the innately indeterministic quantum structure of the world. Or it is a side-effect of the brain's information entropy attempting to "slough off" its waste processing in the course of formulating a determinate conclusion. The latter, incidentally, directly parallels C.S Pierce's notion of belief fixation in which doubt resolves into belief through the process of thought.

More is to be said about how the laws of physics rigidly constrain the knowable. next we will turn to several key constants and physical edge cases, the speed of light, light cones, and black holes, for insight into the nature of the unknown.

This brings us to a critical point. Does physics tell us what can't be known?

In support of this idea of the substantiality of ignorance several concepts from physics have been marshaled. The Heisenberg Uncertainty Principle perhaps the most essential, though it carries its detractors, such as staunch believers in determinism who insist that "god does not play dice." But if there's any truth to

physicist John Archibald Wheeler's claim about the ontological root of physics in information, *It from Bit*, then the search for the foundations of epistemology can begin with physics.

Information unites every natural phenomenon, from atoms to brains. knowledge is physical. As physicist Seth Lloyd has said, "everything in the universe is made of bits. Not chunks of stuff, but chunks of information — ones and zeros. … Atoms and electrons are bits. Atomic collisions are "ops." Machine language is the laws of physics. The universe is a quantum computer."([https://www.wired.com/2006/03/keeping-it-reel/?pg=4]) "Made" here is a strong word —it denotes strict ontology. The universe is not merely describable in bits, it is bits. Knowledge is made up of the same stuff, hence the profound opportunity to link physics and mind. One can pick at this informational ontology and find holes in it, claim that it mistakes the map for the territory. One thing that can't be denied is that it has led to profound discoveries in physics and engineering such as the development of quantum computers.

More insightful concepts can be called on in pursuit of a physical basis for the limits of knowledge and objective ignorance: black holes, light cones, the speed of light, and dark matter. These phenomena share the same feature of organizing the boundaries within which physical information is theoretically accessible and outside of which it is strictly inaccessible. That is, they define what is knowable and unknowable. They provide objective coordinates by which one can demarcate the ultimate physical limitations of knowledge.

I do not claim complete knowledge of these physics concepts, partly because of my own lack of education but also in large part because not everything is known about them even to those at the pinnacles of expertise. This is no real limitation however. Based on their very nature we should *expect* ignorance and lack of full understanding to emanate from these edge cases residing at the border of human comprehension.

Black holes are probably the most explicit example of physically determined ignorance. Black holes are localities of spacetime where the mass has become so densely concentrated that it curves space infinitely. The resultant deformation creates a gravitational sink so deep that not even light can escape it, hence the

perfect darkness of these strange cosmic entities. A darkness which, is so "loud" we are able to detect them by noting abysses of spacetime that are so absent of light that it contrasts sharply and visibly with the surrounding region of spacetime.

Incidentally, a supermassive black hole is at the center of our own galaxy, and there is likely to be one in the center of most galaxies ([https://science.nasa.gov/astrophysics/focus-areas/black-holes#]). We are literally funneling into a black hole as we speak, flushing down the cosmic drain. There is a poetic metaphor to this image. Just as ignorance is said to "drive us" into it, so do inscrutable black holes pull us into the abyss of unknowability.

Black holes can vary greatly in immensity but the largest, called supermassive black holes, are important structural fixtures of the universe, lending it much of its shape. Black holes like much of recent physics are interesting from an epistemological perspective because they are typically thought of in information-theoretic terms. Wheeler, of It from Bit fame, was a world-renowned black hole physicist, and has indeed been credited with popularizing the phrase ([https://www.nytimes.com/2008/04/14/science/14wheeler.html]). Black hole physics reads as if what the physicists who study it are trying to understand is not just a bizarre edge case of general relativity but the frontier of intelligibility itself.

Beyond this limit of intelligibility , the mystery of black holes threatens to upend the entire consistency of physics. The law of conservation of energy states effectively that energy cannot be created or destroyed, and so too information. Yet according to the black hole paradox, first proposed by Stevin Hawking black holes seemingly destroy information, violating the principle of unitarity which states that the present always conserves information about the past. While recent research suggests that black holes may in fact conserve the information they absorb (Donald Marolf, Henry Maxfield 2020 [https://arxiv.org/pdf/2010.06602.pdf]) regardless of these technical findings, it is clear they behave in ways that are odds with the familiar world of regular gravitational non-extremes. In a sense they defy reality as we know it. One cannot simply analyze a black hole like one might another interstellar body and infer its properties. Only through advanced and mind-boggling theoretical mathematics can we hope to

glean any details about their nature.

By virtue of the informatics-based concepts that physicists use to analyze them, black holes are inherently epistemological phenomena. Black holes are classifiable within the threefold ignorance-knowledge relations outlined earlier. They are the ultimate known-unknowns. Black holes are known-unknowns writ supermassively large.

Because the information content inside a black hole's "hidden region" is strictly inaccessible and possibly—according to longstanding arguments— even gone for good, only by analyzing what comes into contact with its edge —called its event horizon— can anything about it be known. The circumference of the hole is easily enough gleaned by relating the diameter to the radius and drawing a line across it from a bright spot at its perimeter to where the darkness stops and the visible universe begins again. Its mass can also be similarly calculated by measuring the so called accretion disk of all the mass moving towards the black hole. This gives you a sense of the intensity of its pull. Therefore only by bracketing the black hole within a boundary of intelligibility can anything be indirectly inferred about its properties. This is similar too how we often measure the extent of our ignorance in general by noting its contrast with the extent of our knowledge.

This alone is a lesson in epistemology. The unknown is richly contextualized by everything that is known. It's possible to know more about what we don't know by reference to the surrounding network of knowledge entwined with it. We know precisely what to rule out about what we don't know about, constructed these bracketed, enclosed unknowns of a special, clear cut character. The situation can be likened to completing a jigsaw puzzle for all but one missing piece. You do not have to look at the piece itself, the positive knowledge, to know where it fits. You can reverse the operation and look at the spot where it's missing to know where it goes.

A black hole's hidden region is thus as strange a physical as it is an epistemological concept. It is a box where information goes and does not come out or at least, where it becomes drastically harder to recover or comes out in a less organized form, as in the so called Hawking radiation phenomenon. If you

had in mind to lock something up and throw away the key, and really wanted to get rid of that key, tossing it into a black hole so that it fell into the hidden region would be the ultimate form of irreversible encryption. Not even the light you would need to become aware of the key would be able to reach you to inform you of its whereabouts.

The mention of light brings up an important tie-in. The speed of light is the upper limit at which information can be transmitted. The speed of light determines the causal structure and order of physical reality. Effectively, nothing can happen faster than the light which transmits information about its having happened, and light conveying evidence of an event must always arrive before an observation that the event occurred can be made. Because light speed and information transmission are so closely tied, any statement about information physics is ultimately a statement about epistemology to the extent that you can't have knowledge without information.

Just as light is the limit of the knowable in physics, so is to see to know. To *have a vision* figuratively is to project a more ideal arrangement on things, to know better. So it's shocking when when black holes are described as swallowing up all light. Metaphorically this is like saying they devour knowledge. So they are cosmic blindspots, abrupt points of total oblivion at which the observable universe cuts off. Not only that but they are sources which objectively *increase* the total of ignorance in the world. They remove data by sucking up the surrounding matter that falls into their maws. Their enigmatic nature is not incidental but a consequence of their destruction or occlusion of of information. They are not only mysterious but they create a mystery by devouring cosmic information.

The light cone of an object specifies the extent to which any two events can be causally related. A light cone is the set of all rays of light that could reach a given object at a given time. As such it defines the causal structure of an event. Anything that has occurred so far from an event that light could not have reached it can have no causal influence on said event as no information would be have reached it in time. Thus only events that fall within an observer's past light cone can be knowable to it, and only events within the future light cone of that

even could possibly be influenced by it. Outside of strange cases of quantum entanglement in which two particles can exchange information regardless of spatiotemporal constraints, this is a way of saying that no two events can influence each other unless there is a particle of light that could travel between them by the time the action initiated. It's important to note that it's not that events outside one's light cone can't be seen, it's that no information could arrive before light. Thus the events outside the radius of the light cone are strictly unknowable because it is impossible for one event to be caused to know about the other event ([https://hepweb.ucsd.edu/ph110b/110b_notes/node46.html]).

To reiterate, a light cone is the collection of all the rays of light that project to and from an event at a point in time. Anything that happened outside the radius defined by this hyperboloid of rays cannot have been part of the causal history of the object. Inversely, nothing outside the scope of the light rays emanating from an object can be involved in its future. Therefore, anything outside our light cone is not directly inferable. Taking the light cone of every event in the universe then defines the causal structure of the universe: we get the set of all potential causal relations and find out what can and cannot become known to whatever else by observation. Without overlapping light cones, two events are effectively causally isolated from each other. There is no way for one to learn about the other. The bare minimum requirement for any empirical knowledge of an event is the conveyance of information by observation, starting with the transmission of light as the upper limit. By appreciating how light cones circumscribe causality we can see how absolute ignorance is a fundamental consequence of physics.

When all is said and done this is point about light cones is frankly academic. For the overwhelming majority of practical cases nearly all events that we might want to learn about fall within our line cones. But the concept of a physical light cone implies an epistemic light cone, a horizon that posits an absolute limit on direct knowability. There is an innate cap on what what can be known about the world built into the nature of physical reality. If our goal is a thorough definition of what is unknowable, then we must factor in light cones.

To know always also means to be caused by. knowledge is caused by the

existence of the fact and the acquaintance with it thereby affected. Without some chain of causality linking events there can be no direct knowledge of it. Information is carried on the backs of a causal trace which fly atop the wings of light. An alternative definition of information as a difference that makes a difference is related to this idea of causality. To make a difference here is to attribute causal powers to a knowable entity such that it is distinguishable from its background and becomes a signal apart from the noise and so is able to cause itself to be known to some observer.

To know is often also to *find out*. If there are facts about how information flow in the physical universe is confined by the speed of light then those facts are also about the upper bounds on the availability of knowledge. Physics tells us that there are limits as to what can be found out. While admittedly a somewhat technical detail, this strongly supports the claim that there is an objective structure to ignorance and that knowledge is restricted to certain provisional areas of informational accessibility.

An objection might be raised at this point. How much information do you really need to know something? After all if we are willing to concede that at least some inductive inferences are correct, valid knowledge can be obtained with incomplete information. In the right circumstances a very small amount of information can go a long way if leveraged by powerful reasoning. They key is that there is *some* information and the *right* information. To have no information is to have nothing, to have no rational justification for claiming empirical knowledge. Not only does the thought that consists of knowledge *use* information but it *is* information. Knowledge, as a feature of the physical system the brain, therefore both is about and consists of information. It is part of a kind of continuous field of information encoded in the structure of physics. The best we can do in the absence of relevant information is “fill in the blanks” using the information we do have, to make an inductive leap and generalize to the unknown.

Indeed, there is a concept in physics called the Bekenstein bound which defines the upper limit of information entropy for a finite region of space. That is, it defines how much information you would need to fully characterize that

system and how much information you can fit in it given the finite energy of the system. For a finite physical system there must be a finite amount of information to define it. The Bekenstein bound means that there is a limit to how much information you can store in a region of space with finite energy and matter. That would imply then, that if the amount of information in the universe is greater than the amount of information you can store in a brain, then the brain is constrained according to the laws of physics to a fixed upper bound of structural ignorance. In the same way there is an upper limit to how much information can be packed into a microchip before there are no more degrees of freedom at the quantum scale available to encode more data. Consequently, black holes are explicable in purely informational terms in this way. The more you compress information into a finite region of space than the bound allows, the more you hide that information behind a black hole (Three Roads To Quantum Gravity, Smolin, 2002).

The difference between deduction and induction in terms of information accessibility can be delineated by the difference between having *contiguous* versus *fragmentary* information. With contiguous information access, every element of the computation is linked to another adjacent to it; everything is provided up front and closed under its own rules. Fragmentary information access is spotty. Localized quadrants of information can be broken up by gaps which require the present information to be leveraged over the gaps. The necessity for the amount of information characterizing a region of space to be discrete has led physicists to the conclusion that space must be granular rather than continuous at the smallest scales. Otherwise an infinite amount of information would be needed to specify a region of space ([https://arxiv.org/abs/1203.6191]). But if a lawful parity exists between fundamental physical quantities and information bits, nowhere can it be true that an infinite amount of information be required to describe an otherwise finite system if the Bekenstein bound is to hold water. This is another way of saying that knowledge is limited by the limitations of space itself.

We may not need to know every piece of information in a contiguous information system to reason about it accurately. What matters is that when we

go to confirm the argument, the information we presently lack is discoverable. For fragmentary information systems, only one part of the object we're picturing with our argument survives information decay. By this token, the universe taken as a whole is a fragmentary information system, punctured by black holes and carved up into partially discoverable regions by light cones. Black holes have taken up some information and will never give it back, at least not in usable form, while light cones and the speed of light fracture the continuity of information flow. For most practical purposes however, the region of the universe we're interested in is an informationally complete system at least in principle. Yet the fact it is subject to entropic dissipation means that it may not be able to access all that information, to say nothing of the many other practical reasons that knowledge may be inaccessible to us.

Now to take stock. The general tendency in modern physics toward digitalization is of deep epistemological interest. It can be taken to express an equality or isomorphism between physical quantities and information quantities. This would mean that the basic objects of physics and the mind are expressible in the same language of informatics. The essential equivalence between mind and matter is to be found in them being equally subject to the laws of computation. knowledge and information are not reversibly identical in the sense you can have information without knowledge but you can't have knowledge without information. Knowledge is made out of information (a particular brain state) and is about information. Information is a characteristic of the state of a system, it is inherent to how a thing relates to itself or to another. When physicists such as Seth Lloyd speak of the universe as performing computations and the laws of physics as a kind of programming syntax, it should present itself as a great opportunity to blend mind and matter, thought and substance.

The uniformity of nature under information, brains and black holes included, means that certain statements about physics are also statements about the mind. Of these we have noted the uncertainty principle, black hole physics, the speed of light, light cones, entropy, and so on. If mind and physics are united by the same descriptive language then certain riddles of physics such as the Copenhagen interpretation vanish. The observer and the observed are not

mysteriously divided, disconnected entities. They are part of the same unified information system. We are all part of a "participatory universe" to use another of Wheeler's phrases, subject and object are united on a fundamental level. Everything that applies to one applies to the other, where previously, we used to consider material and mental phenomena as descriptively incompatible. There was no language to describe both of them in the same terms. Information theory provides just that.

The state of eternal epistemic ambiguity toward the world at large in which the human being is permanently embroiled is a fundamental result of the structure of nature. Curiosity and the drive to discover through exploratory behavior is as much a response to the laws of nature as hunger is a response to the laws of energy and the need for an external source of bodily fuel. Information entropy reduction is an essential feature of thinking things.

The ultimate bounds of ignorance are reflected in these loopholes and edge cases embedded in the fabric of physical reality. That nature places absolute limits on knowledge is a strong statement, and one that goes beyond physical limitations and into the computation as we will soon learn. But this is not to inspire pessimism. By investigating the intricate ways in which we are doomed to ignorance, by appreciating the true meaning of our restrictedness, we can effectively mitigate the perpetual lack of incomplete knowledge and its consequences. We can have a more or less complete understanding of our epistemic incompleteness.

Unknowns aren't simply always detrimental to theories, they can be semantic components of them. Consider hidden variable theories. A hidden variable theory posits unobservable measures underlying quantum (or other) phenomena that balances the equations so that quantum indeterminacy is banished. Such theories present an interesting attempt to explain observable events in terms of factors which are not knowable, or at least not knowable with the transparency of observables.

It's a bit of a denouement to point out that hidden variable theories have more

or less been discounted as a viable approach to the mysteries of quantum mechanics, with the possible exception of what are called local hidden variable theories. (Roger Colbeck; Renato Renner (2011). "No extension of quantum theory can have improved predictive power". Nature Communications. 2 (8): 411. arXiv:1005.5173. Bibcode:2011NatCo...2E.411C. doi:10.1038/ncomms1416. PMC 3265370. PMID 21811240.) Nevertheless, it's the idea that unobservable facts can influence events itself which prompts some intriguing discursion.

Can what is unobservable be truly knowable or is it necessary that what is knowable be observable? Perhaps the answer to this question depends on how richly you define observable. Observation is not necessarily just perceptual detection by the bare senses. It can include readings from fine-tuned instruments that register humanly imperceptible events, and it arguably includes any form of exact measurement that defines a definite quantity associated with an entity which can be verified by experiment or some other precise method. Observation does not even require the object of interest to even be specifically observed, rather, by noting how the causally influences of other more readily observable events affect it its existence can be inferred. Particle colliders for instance smash two highly concentrated beams of particles and then capture snapshots of the resulting scatter pattern of ephemeral subatomic and sub-subatomic particles released from the impact. From these bursts of baryonic nuclear radiation scientists are able to infer information about the enigmatic properties of dark matter which is otherwise virtually undetectable([https://scitechdaily.com/new-antimatter-experiment-at-large-hadron-collider-will-help-with-the-search-for-dark-matter/]).

Certainly there are truths which are more abstract or restrictive in nature which are not immediately subject to observation. So called black boxes in engineering are modular systems the inner functioning of which is encapsulated and inaccessible to the other mechanisms. Their hidden functioning is certainly causally effective and real though it cannot be observed. The difference between a black box and a hidden variable theory is that in principle one could crack open a black box and analyze its internal operations. It is only designed to be

unobservable to the other systems because it is expedient and simplifying to restrict access to just its inputs and outputs.

Similarly nobody observes an imaginary number in the typical sense that they can be counted like apples, although they can employ it in symbolic representations and demonstrate their mathematical validity and utility. One might say that mathematicians still observe mathematical truths by deriving them through perceptible notation systems and by manipulating mathematical laws as inexorable and constraining in the world of math as gravity is on the physical. Such truths can also be inferred analytic a priori, through the logic of the terms itself, without having to relate or synthetically attach them to some external evidence to verify.

Though hidden variable theories are more similar to the case of imaginary numbers than black boxes in that they are purely mathematical constructs, they require stronger truth conditions: they must be both mathematical and physically true. Facts which are strictly unobservable are indisputably more ambiguous than those which can be confirmed to exist through some concrete act of observation. What else is unobservable? Everything that is imagined and fancied, leprechauns and Dracula: everything that does not objectively exist. Thus unobservable phenomena have at least as much in common with the fantastical as they do observable facts. They split the difference by being unobservable like the purely imaginary and by being facts, supposing they are true. Is it possible to confirm the truth of unobservable facts, or do all means of knowing their truth rationally fall short?

Again even analytic a priori truths are arguably subject to *internal* observation, given that the proposition is inspectable and its meaning derives from its intrinsic logic. There is no need for external verification because of the tautological essence of such propositions. Even though “all bachelors are unmarried” could be externally confirmed by interviewing every bachelor and verifying the fact, it is unnecessary. So on this interpretation hidden variables can be ruled out as not being a priori, they are by no means logically nor physically necessary; after all quantum physics works just as well without it. Both synthetically and analytically therefore, hidden variables aren’t knowable. What

means of verification is still available?

What's most interesting about hidden variables is the fact they are hidden: some ignorance of them is presumed. What is unobservable may not be strictly unknowable but it is certainly less knowable than something directly inspectable. These variables are summoned into physics to rescue it from the intolerable possibility of objective indeterminacy. They are complementary counterweights to perceived incompletenesses in quantum theory. Such variables are a paradigm of *structural* ignorance. Perhaps the only way that they are knowable is that it is possible to formulate a variant of quantum theory which banishes uncertainty using it—something about the mathematics used in hidden variable theories fits the experimentally confirmed structure of quantum theory. This creates a more harmonious picture of physics, although beauty does not necessarily equal truth. Hidden variables, as unobservable assumptions, are thus concatenated with known truths, they are known unknowns and perform a structuring role in securing a purely deterministic model of physics. Regardless of whether such theories are accurate, it's an unusual and compelling idea from the perspective of ignorance theory.

Having now gone into some detail about the physical constraints on knowability, the next step is to get into computability. Computation and its limits is, after physics, the most root level constraint on knowledge, to the extent that the mind justifiably can be regarded as computational in nature and that all reasoning processes are rule-based.

Part II Sectiion II: Logical and Computational Limits

The computational theory of mind has become something of the central dogma of cognitive science, as close to a unifying explanatory framework as it presently has. As the name suggests the theory proposes that the mind is fundamentally computational in nature, that thinking *is* computation, symbol manipulation. It is not just that the mind manipulates language or symbols as we are familiar with them such as spoken or written words or various notations.

Rather on a deep embedded root level of its operation, there is what Fodor called a language of thought, that concepts and ideas were ultimately tokenized symbols generated by fixed rules of mental representation.

Early hints of the idea can be traced back hundreds of years to Leibniz' dream of a*characteristica universalis*, a universal alphabet of thought that would enfold all ideation into a single symbolic formalism. Leibniz' pursuit of a universal language of thought never materialized, but was the first intimation of the notion that such a codification of the mind's rules and contents might be conceivable .

Leibniz' potentially fertile idea was to lay fallow for centuries, more or less unappreciated or considered a historical curiosity. With the development of both brain and computer science in the middle of the 20th century it was to seem that the time was ripe for the idea to be resurrected, even if unintentionally and in a form that reversed Leibniz' assumptions. In 1943 neurophysiologist Warren S. McCulloch and logician Walter Pitts teamed up and published a paper titled *A Logical Calculus of the Ideas Immanent in Nervous Activity.* ([https://www.cs.cmu.edu/\~./epxing/Class/10715/reading/McCulloch.and.Pitts.pdf]) In this paper they observed that because the states of neurons are boolean valued all-or-nothing (true or false) this made them amenable to propositional logic. These action potentials of neurons were also vectors, they contained a direction, in terms of their outgoing afferent connections to its associated signal-recipient neurons, and a magnitude, the electrochemical signal it outputted. The researches proceed to develop a calculus to codify the implications of this idea, which would comprise a major contribution to the future development of neural network artificial intelligence models and influence the language of thought hypothesis behind the computational theory of mind.

This notion of a calculus for brain activity is not a natural extension of Leibniz' universal concept language, but more of an inversion of it. Because it is taken to mean that such a language of thought is not descriptive of thought but ontological—it is not that there is a logic to describe the organization of mental representations. It's that the order of concepts *is* this logic and thought is and is caused by the implementation and execution of its rules within neural tissues.

The goal is not to devise a universal language to express all concepts, but to equate all concepts with such a language. The mind and its operations therefore is not just describable as a formal system, it is a formal system. Collections of interlinked neurons form circuits that are on a computational level indiscernible from a formal logic consisting of a vast set of propositions composed of binary valued terms. At the brain works, such propositions are stringed together in combinations that can be generated and varied indefinitely.

Toward the end of their paper McCulloch and Pitts make an intriguing remark about what their model implies about the nature of ignorance as it is instantiated in the computational brain:

Specification of the nervous net provides the law of necessary connection whereby one can compute from the description of any state that of the succeeding state, but the inclusion of disjunctive relations prevents complete determination of the one before. Moreover, the regenerative activity of constituent circles renders reference indefinite as to time past. Thus our knowledge of the world, including ourselves, is incomplete as to space and indefinite as to time. This ignorance, implicit in all our brains, is the counterpart of the abstraction which renders our knowledge useful. The role of brains in determining the epistemic relations of our theories to our observations and of these to the facts is all too clear, for it is apparent that every idea and every sensation is realized by activity within that net, and by no such activity are the actual afferents fully determined. (p. 13)

What this passage says is that while operations of the brain are closed under such neural computation, the stochastic, semi-regular disjunctive processes of such activity means that the previous state cannot be recovered from current state. The fact that the state took one path of disjunctive switches at one time is not predictive of whether it will ever reenter that state. Both spatially, in terms of which structural pathways the circuit propagated at each loop cycle, but also in time, as the "regenerative" process iterates without stopping and therefore without differentiating its steps into independently parseable static

configurations. No state of the net can be explained in isolation apart from the entire irrecoverable, "blurred" history of the circuit's past activity. Recovery of the past in space, in terms of retracing the disjunctive steps that the system took as it switch-gated one way rather than another, and recovery in time, as each state is blurred into the continuous cycling of the circuit, means information about its state is ephemeral and soon lost. From this fact of spatiotemporal information loss in the brain, the authors proceed to make a statement which is startlingly and profoundly relevant to the theory of ignorance:

> There is no theory we may hold and no observation we can make that will retain so much as its old defective reference to the facts if the net be altered… Thus empiry confirms that if our nets are undefined, our facts are undefined, and to the "real" we can attribute not so much as one quality or "form." *With determination of the net, the unknowable object of knowledge, the "thing in itself," ceases to be unknowable. *\[emphasis added] (p. 113)

Thus the unknown known, the deep neural encodings that underlie knowledge, is implicitly imprinted in the logical structures defined by the specification of the net. Each previous state of the net, accordingly, is an outdated model of the facts given that the net is alterable and ever adjusting. Our knowledge is always out of date because the net never ceases its operation and arrives at a completed computation, it is always fluid. The underpinnings of knowledge in this case is identical to the unknowable object of knowledge, thus we come full circle and see that the hidden substrate of that knowledge is itself knowable as neural logic. Though in principle the unknowable known as it were, ceases to be such when considered that it has a definite structure encoded in brain circuits, in practice it remains unknowable. What is telling about the authors' philosophical remark is that the brain is indeed the object of ignorance and its operations are ultimately transient knowledge states that briefly surface out of ignorance as the system's configuration is never preserved in space or time. This is quite in line with what we have said about unconscious knowledge and the vagueness of the knower.

Decades later McCulloch and Pitt's preliminary research was to be vindicated and extended by the development of artificial deep learning neural networks

which solidified the fact that brain-like and brain-inspired processes can be schematized and implemented as computational models that achieve real results in perception, reasoning, pattern detection, chess playing, natural language processing and other marvelous feats of braininess. There are quite a few places one can take an interpretation of these results. Working with the premise that the mind may indeed be a formal system, we shall see if what is known about the constraints of said systems must therefore also apply to minds. The putative necessary connection between formal and brain logics that suggests they are intersubstitutable *salva veritate* is by no means a given, but it is nevertheless a bounteous and fruitful avenue of investigation for reasoning about unknowability.

The Limits of Computability and Cognitive Closure

Consider the moment when you realize you don't know something. That moment when the fluid flow of one's understanding halts, and the indefinable shadow of one's ignorance appears before you as a vague *presence of absence* The natural response to the unknown is to make it known. Through some process of transformation the unknown becomes known, a nothing becomes a something, and the algorithms of knowledge growth continue to traverse over the space of all possible knowables. At first glance the unknown fact appears to be "out there" yet occluded; what we come to know was there all along. But what if there were certain cases where that unknown cannot be made known? What if there is no pushing further, and the limit of our understanding that we realize marks out a final boundary? This brings us to the problem of unknowability.

Some things can't be known only relatively to the agent's positioning in the scheme of the world, even if they might be knowable in principle. Consider arriving at the station to learn that the train has been delayed. One mile outside the station, a tree fell on the tracks and blocked the train. One couldn't but act under the assumption that the train would arrive at the stated time, and to

assume otherwise would go against the most expected outcome and cause you to miss the ride. From the agent's perspective no immediate information source is available to inform them of the delay, although the concept that a train can be delayed and that this train in particular will be is entirely understandable. This relative unknown is only unknown because what is missing is a series of information links connecting the fact of the delay to the agent's awareness. Context here constrains the knowable. This is a kind of ignorance that is *forced*, that is, for which one cannot be held responsible. Nothing could have allowed you to anticipate it without making what at the time would seem like an improbable, unreasonable decision.

Perhaps some things simply cannot be known in principle. If this is true, then the next best thing is to know why we can't know certain things. The reason some things cannot be known categorically is much deeper than the situational issues of relative ignorance. This absolute ignorance is a side-effect of the intrinsic nature of reality, the structure of our minds, and perhaps even the logic of computation that underwrites the symbolic order of thoughts. Contingent, relative ignorance is merely something you missed, or rather something that missed you. This deeper form of ignorance however, is lawful. Just as the laws of nature help us to generalize over the space of facts, so do these laws, but instead of telling us what we can know in generic terms, they tell us what we can't.

There is a distinction between contingent and systematic knowability. The generality of the laws of physics lets us know the causal form of general facts. Contingent facts are harder to ascertain because of their transient, one-shot, high specific character. They are time and context dependent. They stand for a time and then dissolve into the flow events. Systematic facts are those which structurally undergird a class of phenomena in all cases; each time an object falls to the ground is an opportunity to relate the event to the law of gravity. But there may also be some absolute facts which do the opposite, they tell us what can't be discovered.

Underlying the problem of unknowability is the question of *cognitive closure*. Are there some truths that are indiscernible to the human mind by virtue of its very constitution? Are we only capable of thinking and knowing

about a relatively constrained local area of all possible thoughts and knowledge because of the construction of our brains? It has been argued for example that the hard problem of consciousness is cognitively closed. The brain, perhaps, is evolved to recognize only certain internal and external states through the senses, but it is not built to monitor this same perceptual apparatus using the same framework it applies to everything else. The issue is similar to the proverbial eye which cannot not see itself within its field of vision. There are certain facts which the architecture simply does not support. This is somewhat after a fashion of Kant's Categories of the Understanding: we can know Quality, Quantity, Relation, and Modality, but these also constrain what we are able to think about. Facts without these forms are unthinkable, because all thoughts must conform to this format.

The question of cognitive closure is perhaps best approached from a theory of the limits of computation. The question may very well, as we will soon demonstrate, belong to a class of undecidable propositions which are characterized best in terms of their formal, computational, and logical properties. How can one prove such a position? One would need a means to show that the question of cognitive closure is itself not closed to us, that it is a relative unknown rather than an absolute unknown. As Wittgenstein writes in the introduction to the *Tractatus*: "Thus the aim
the book is to draw a limit to thought, or rather--not to thought,
but to the expression of thoughts: for in order to be able
to draw a limit to thought, we should have to find both sides of
the limit thinkable (i.e. we should have to be able to think what cannot
be thought)." In other words, we cannot define the unthinkable because that would mean we could think it. There is never an exact demarcation between what might be unknowable in principle and what you yourself simply cannot fathom. Hence Denett's response to those "mysterians" who claim consciousness is inexplicable: "try harder."

Before pursuing the claim that cognitive closure is provable, consider the counterclaim that the mind is universal. In his book *The Beginning of Infinity* physicist and quantum computing pioneer David Deutsch refutes the idea of

cognitive closure through the concept of *computational universality*. To illustrate the concept Deutsch compares the earliest pictographic writing scripts with the historically newer alphabetic scripts. Pictographic scripts are inefficient in the sense that they must have a single pictogram for each word-concept. The total data structure of the script then becomes a big list containing each symbol for each concept. The list must be rather clumsily updated whenever a new symbol is needed for a new concept. The "reach" of such a script, as Deutsch calls it, is bounded. It cannot generate every possibility using a set of elementary units. Alphabetic scripts in contrast operate on a lower level of meaning, the phonetic. It encodes the basic units of sound which when combined produce any possible grammatically valid word in the language. In result alphabetic scripts are more efficiently extensible and compact. A small set of a few dozen atomic symbols can be recombined to yield a theoretically infinite number of new words to accomodate an indefinite amount of new concepts. Such an arbitrarily extensible system is universal. Its reach is unbounded.

In the realm of computing Turing completeness is the digital corollary of universality. A Turing-complete computer can run any program on it disregarding practical constraints such as how much memory it would require or how much time it would take to execute. Following this line of thinking, according to Deutsch people are "universal explainers" able to combinatorially rearrange elementary concepts to specify a causal theory describing any pattern. Humans have limitless explanatory reach.

The concepts of universality and Turing-completeness would seem to challenge the concept of cognitive closure, suggesting that while there may be indefinitely many unknown unknowns there are no unknowable unknowns. Any system that acquires universal reach has unconstrained generativity. It's never going to be stuck with a static list of fixed elements. It can always generate more combinations and accomodate new additions. The catch here is that these are discrete symbol systems must be subject to whatever equally universal computational limitations apply to such symbolic processing.

Whether thought is intrinsically symbolic is a source of extensive debate in the cognitive science literature. Some, like Fodor, propose a language of thought

hypothesis, or "mentalese" that has its roots in Leibniz' pursuit of a universal "alphabet of human thought." " The alphabet of human thoughts is a catalog of primitive concepts, that is, of those things that we cannot reduce to any clearer definitions" wrote Leibniz in the treatise on the subject, *De Alphabeto Cogitationum Humanarum*. There is in these treatments a strictly computationalist interpretation to the cognitive process, hence the very popular computational theory of mind. To think, in effect, is to unconsciously write out statements in a subconscious code. Accordingly, one might expect that any definable limits to computation in general must extend to mental computation in particular.

The flaw with such a view of universally expressive languages is that its ability to express all possible statements within its reach does not mean it can express *everything*. The question then becomes if *all truths at large* are expressible. Whether, that is, there is at least one truth that is inexpressible. Universality is to be found in a system which is capable of generating all the statements of every non-universal system (completeness). A desktop computer is universal, and thus it can run a calculator program on it, but a calculator is not universal, so it cannot run a PC. This does not strictly entail that everything can be expressed within a universal symbolism which can only produce statements with reference to itself. No ordinary language sentence can express what higher math expresses, even though a natural language with universal grammar can express everything that can be expressed in that language. The fact this is so is why math was invented in the first place; natural language was insufficient. If thought itself is a kind of language, then Wittgenstein's search for the limit to the expression of thoughts is not so far from a search for the limit to thought.

Enter Gödel's incompleteness theorems. According to the incompleteness theorems, there are absolute limits on the decidability of propositions within any axiomatic logical formalism. The theorems prove there are statements expressible within any consistent formalism which are true but not provable within that system. Furthermore, and more crucially for our purposes, according to the second theorem a consistent axiomatic system cannot prove that it is itself consistent. As Davies writes, "the importance of the Gödel theorems is that by

mixing subject and object, it demonstrates how even at the fundamental level of logical analysis self-reference can produce either paradox or indecision." (God and the New Physics). If we are to preserve a uniformity of information across nature, then these considerations about the limits of computability must also be taken as statements of the limit of knowledge.

A brief proviso before proceeding further. The reader can find many more technically sophisticated analyses of the incompleteness theorems elsewhere. The intent here is not to expand on the technicalities of the theorems, but to offer a novel interpretation about their significance within the framework espoused in this work. Attempts to generalize the theorems to areas outside of mathematics are not new, but the case can be made that there is indeed such an insightful generalization to be made, and that it can be advanced in such a way that provides the most accurate and fruitful presentation of its meaning.

To understand the logical meaning of the incompleteness theorems in the clearest light, it's best to look at them in comparison to the concept of computational universality mentioned previously. Other terms for computational universality are *recursively axiomatizable* or *recursively enumerable*. In the first case, a formalism is considered complete if there is always a mechanical algorithm that permits one to decide if a give statement within it is an axiom. Recursive enumerability is the rule of set theories that means you are able to derive a list of all the numbers in the set. It was said previously that just because a system is universal does not mean it has unlimited applicability. Indeed, the incompleteness theorems show that what may be true but unprovable in one formalism still yet may be provable in another. It does show however that any consistent axiomatic system has bounded truth-proving capabilities. It has truths which it cannot prove to be true, and it cannot prove its own consistency with reference only to its internal inference rules. It is limited in both axiomatizability and enumerability.

The implications of these results are as follows. An inconsistent formal system is trivially complete, because anything is provable within an inconsistent system. There are no logical restrictions on the evaluation of the truth conditions of a

statement. If both P and not P, anything goes. So to be substantial an axiomatic system must be consistent. If it is consistent, however, then there are statements within that formalism which cannot be decided as both true and provable. Furthermore, the system is unable to demonstrate its own consistency, which on a level deeper makes it self-referentially and recursively inconsistent (incomplete).

The essential mathematical meaning of the incompleteness theorems is that mathematical knowledge is imperfect. There can be no perfect, universal axiomatic system for which the highest standards of mathematical truth can be obtained. It is impossible, therefore, to pursue mathematician David Hilbert's foundationalist dream to ground all mathematics in an elementary arithmetic or set theory, for example.

To see the wider ramifications of these results, it's helpful to ask what the incompleteness theorems mean for the human mind, for computing, and most broadly, the universe itself. The material equivalence can be rendered as such:

There are statements within any axiomatic system that are true and not provable and any consistent axiomatic system cannot demonstrate its own consistency <——> *There are truths within any mind that would constitute valid knowledge but are not thinkable, and there are some truths which are undecidable.*

Which is to say, if the most powerful paradigm of precise reasoning, mathematics, is fundamentally limited in its truth-capabilities, it's natural to ask if the mind which reasons with such a tool is similarly limited. Loosely speaking, a tool expresses the horizons of the toolmaker or tool user; the extent to which it can invent and use its inventions expresses the extent of the user's powers. By this token we can take recursive computable functions as a stand-in for cognitions which terminate at definite factual knowledge.

Indeed, an implication of the Church-Turing thesis its that there are countably many computable functions, but uncountably many uncomputable functions. Simply, this means that only those numbers which terminate at a definite, non-infinite value are computable, and there is a vast space of functions which do not terminate at finite numbers which cannot be represented in a computer.

According to the Cantor diagonalization method, infinitely infinite numbers can be derived, which cannot be mapped to an infinite series of natural numbers. If you were to draw out a set of numbers

\{000000….,

1111111…,

1010101…, …]

and then draw a line diagonally through them, the number you get by crossing the set cannot appear in any of the set's members. You can draw an infinite number of such lines, because each number in the set is infinite, meaning, strangely enough, there are infinite varieties of infinities.

If to know is to compute, then there is a vast space of unknowable, uncomputable numbers. Each of these results falls within a particular category of known unknowns, or more precisely, known unknowables. They are methods of establishing fixed epistemic horizons. The question is if they can be taken out of the technical box in which they reside and generalized to the mind and world at large.

To proceed further, at first the meaning of math in broad terms must be established. What is the ontological status of mathematical objects? Intuitionist? Logicist? Platonist? Formalist? There is no place here for a thorough examination of the options, although in passing it can be said that Gödel is widely believed to have put an end to the logicist program and was a platonist himself. In effect math cannot be reduced to logic. Of all the options, perhaps the most attractive for our purposes is the cognitivist: that is to say, math is derivable from concrete empirical findings in cognitive neuropsychology. In *What Mathematics Is*, Lakoff and Núñez write, "if there are of foundations of mathematics, they are *conceptual foundations* —mind-based foundations. They would consist of a thorough mathematical idea analysis…showing how these concepts are ultimately grounded in bodily experience." Lakoff and Núñez emphasize there is nothing transcendental about math. It is a product of the structural constraints of the body and its neural capacities. This stance opens up a path to basing mathematical truth in empirical data about cognitive architecture.

At the same time math is widely embraced as one of the most precise and

universal of thinking tools, in fact as the paradigm of precise and universal thinking itself. The cognitivist position, however, does not draw a sharp distinction between math and the rest of cognitive processes. There is nothing mysterious about mathematical cognition, which, as the presumptive queen of the sciences, would place it ironically in a strangely exclusive scientifically inexplicable blindspot. An anomaly presents itself: math allows mathematicians to think about extraordinary phenomena that are not comprehensible through usual sensory channels and yet math itself must be explainable through the basal underlying structures of ordinary cognition and perception. Lakoff and Nunez nevertheless argue that fundamental mathematical concepts can be reduced to primary cognitive concepts. For example, arithmetic is reducible to the phenomenon of subitizing—in which even children can perceptually identify small numbers of objects in their gaze, and note when an object is removed (subtracted) from the visual scene or included (addition). Set theory is explainable in terms of container logic, the perceptual intuition that things can fit inside one another. Trigonometry is derivable from the fundamental cognitive construct of a trajector: the sensorimotor intuition that an object can have impulsion that tends to draw it toward a specific predestined direction. And so on.

To reverse the terms here then, a statement about math is a statement about the mind. On this basis, therefore, it would seem the window is open to validly generalize the incompleteness theorems as both a statement about math and a statement about them mind. This is a bold claim, and it must overcome some serious demonstrative hurdles before it can be embraced in good confidence. Are the terms reversible?

Gödel himself took his theorems to mean there was a transcendental property to the human mind which enabled it to ascertain truths which could not be derived through mechanistic, purely algorithmic processes (Gibbs Lecture, 1951). In which case the entire project to mark the limits of the humanly knowable to be in line the limits of the computable falls apart. "I have not shown that there are non-decidable questions for the human mind, but only that there are no machines that can decide all the questions of number theory. The human mind

infinitely surpasses the power of any finite machine. \[emphasis added]" To infinitely surpass is quite something, and it is hard to accept such a position without supposing it requires an unjustified leap of faith. And to surpass *any* finite machine is to assume that there is no conceivable futuristic supercomputer that can ever reach the capabilities of the human mind. The implication is that the human mind is not an elaborate biological implementation of a Turing machine; that at least some behaviors, such as those of intuition, are not algorithmic. But the brain's principles of operation are not constituted as a formal system in the conventional sense in their most schematic outlines. The mind is not a literal formal system. Yet is it really true that the brain's design architecture is not formalizable? Isn't that the goal of artificial intelligence, or even the scientific project of mapping the brain? Wouldn't this be a claim that had to be sounded out empirically? After all at this point we are no longer talking about formalisms, but brains, objects subject to empirical, more than logical, criteria.

Solid support for the idea that mental processes are formalizable is the success of brain-inspired machine learning techniques. Such deep learning algorithms have enabled computers to do things which used to be thought beyond their scope, such as image classification and natural language processing. Suppose that the effectiveness of these techniques demonstrate some underlying unity of principle across the respective organic and inorganic implementations in brains and machines. Though granted these algorithms not a one to one conversion of brain logic to the machine logic, some overarching conceptual homology holds for both. The reasoning for this is quite clear: deep learning works, and it is explicitly inspired by how the brain is believed to function. To the extent that this is true and can be further demonstrated by future technological advances, Gödel's secret sauce of the human mind will continue to be pushed back into an ever shrinking corner. All of this is building up to the idea that Gödel had it precisely backwards: his theorems do not prove humans have a special power over and above machines, but that we are exempt from the same self-referential loopholes that short-circuit consistent formal systems.

Cognitive science has proceeded in leaps and bounds since Gödel's day. The discovery that most thinking is unconscious and that intuitions, sudden eureka

moments, are the tail-end conscious peaks of deep unconscious processing puts the claim that the mind infinitely surpasses algorithms to rest. It's only because there is no awareness of it and because we receive its output in conscious thought from apparently nowhere that it seems like there is no algorithm that can explain intuitions. So the burden of proof is on Gödel. How do we know such supposedly non-algorithmic insights aren't just differently algorithmic? Executed that is, by special neural programs with which the brain comes distinctly equipped?

Perhaps as Penrose (Shadows of the Mind) has theorized, thinking occurs on a deeper level than the neuron and synapse. Microstructurally, the brain has three distinct layers of processing: the neuron action potential, the synaptic interchanges, and microtubules embedded within the cytoskeletons of neurons. This lowest level is non-computational and indeterministic. Penrose sides with Gödel in agreeing that, if there are provable limits to the consistency and completeness of all computational systems, the computational theory of mind is false. The argument in support of this is tortuous and wending, and here will be presented only the most abridged, simplified version. In short, the theory is called orchestrated objective reduction (Orch OR). It asserts that the microtubules engage in a variation of superposition collapse, a spin on the wave function collapse mentioned in the previous chapter. Essentially, the quantum state remains in superposition until the microtubules increase the curvature of spacetime to such a degree that it collapses. This hypothetical phenomenon supposedly is implicated in conscious states. This process is not computable and cannot be specified by a mechanical algorithm. It is, Penrose suggests, the source of Gödel's special sauce.

Some of the refutations that have been proposed against this position are quite alluring from the perspective of the theory of ignorance. Both Minksy (1991) and Chalmers (1995) advance the idea that Penrose is wrong because people, namely mathematicians and logicians, can be wrong. There is in other words, "no basis for assuming that humans are consistent -- nor is there any basic obstacle to making machines use inconsistent forms of reasoning." (Minsky, 1991) This is another way of saying that human thinking and the structures that

underlie it mustn't necessarily rely on the same criteria that apply to formal logical systems, invalidating Gödel's insistence that the theorems apply to both, and that humans inability to conform to it is proof that there are non-algorithmic aspects of cognition. Conversely nor does this rule out the possibility that machines cannot be made that exhibit the messy "loose" forms of reasoning that characterize so much human thinking. One could wager that AI neural networks implement precisely this kind of sloppy logic, and the fact that it is not so able to associate loosely rather than constricted by rigid rules is actually a benefit. They are not exactly correct in their conclusions, unless we deliberately put in place a supervisor that serves as an external truth decider for the system's outputs. But they can detect patterns.

Penrose and by extension Gödel's claims rest on a principle of reflection, "the type of 'seeing' that is involved in a reflection principle requires a mathematical insight that is not the result of the purely algorithmic operations that could be coded into some formal mathematical system." The flaw in their argument is that they lack the ability to reflect on the possibility that brain operations are under no such obligation to be consistent or to suffer the indignities of logical self-reference.

For one thing the premise that the most important bits of human reasoning are, in some sense *irrational*, that is, not specifiable by any concrete algorithmic description, would also fix it such that such processes are unknowable to reason. A scientific explanation is in one view a mechanistic blueprint, an output of a mechanical process. Penrose seeks to reason algorithmically about the unreasonably non-algorithmic. For what it's worth, this gets us to an important point about a cut off point at which our capacity to understand ignorance is sharply delimited. Suppose that consciousness *is* explainable in terms of quantum gravity. Then that would mean that such a process is schematic, it can be specified by its embodiment in the brain as a biocomputer. Either that or *there is some mysterious and inexplicable property that we can't know that allows us to know anything at all*.

This hypothetical property is a kind of linchpin proposition the truth of which regulates the limits of the knowable. It is the point where understanding cannot

understand itself because of some strange self-referential or reflexive aspect of introspection. For it to be otherwise, understanding would have to be like a self-replicating von Neumann universal constructor: thought would have to be able to devise an exact copy of the process that derived it. More than metacognition, or thoughts about thoughts, a thought would have to be about itself, in a demonstration of self-repetition and self-construction.

This last point might be difficult to understand so it bears elaboration. The fact that we find it inconceivable to formulate a thought that tells us precisely what the thought is and how it's made, is the limiting case of the problem of cognitive closure and the same reason behind the seeming insurmountability of the hard problem of consciousness. If consciousness is, according to Husserl, always consciousness of something, that something cannot be consciousness itself. This is implicit in the claim that consciousness must have an object. That object cannot be consciousness itself, because consciousness is directed: it is about something else, rather than "about aboutness".

Things are admittedly getting a bit loopy here, but it is a necessary trouble when dealing with self-referential problems. These paradoxes of self-reference are instrumental in defining the limits of inferential systems. It is like the eye that cannot see itself. Has this discussion brought us any closer to answering whether there are some truths that are strictly unknowable? The incompleteness theorems are statements about axiomatic systems in relation to Peano arithmetic, a set of axioms used in mathematical logic to attempt to provide logical foundations for number theory. Math is widely considered the paradigm of rational thought. The contention then hinges on whether the effect speaks for the cause. Whether that is if statements that are true about the products or process of mathematical thought are also true about the engine (brain or calculator otherwise unspecified) which yielded them. Fittingly in accordance with the first theorem, this puzzle may well be undecidable. Presumably you would want to use math to prove the claim, but that would mean you would have statements that are true but unprovable, leaving you right where you started. This is Wittgenstein's limit. The question of cognitive closure is itself cognitively closed, because we are unable to express conditions that would prove it one way or the other. Logically, however,

if undecidable is taken to mean closed, then the issue of undecidable cognitive closure is itself closed. Then cognitive closure is true at least in the minimal limiting case of its self-application. If it's true, then it's decidable. If it's decidable, then it's false. What is going on here?

The way out of this logical maze is to raise the issue of second-order decidability: that is, of decidably undecidable propositions. The answer to the question if a proposition is decidable or undecidable is yes or no in either case. Because we can answer that cognitive closure is undecidable, we arrive at one solid and firm foundation for knowing well what we can't know. We return to Socrates' dictum: "I know that I know nothing." There is at least 1 item which is cognitively closed to us, the question of cognitive closure. It is the ultimate known unknown, the one that gives us the rest.

Decidably undecidable propositions are statements about problems for which we can say there is no algorithm for determining if its answer is true or false. The ability to decide that a problem is undecidable seems to be what Gödel had in mind when he spoke of the human mind infinitely surpassing machines. Algorithmic behavior in this regard renders either inconsistent (occasionally true or false results) or the program runs forever. But humans seems to have a special power of recognizing the futility of the question rather than to mindlessly go on in an infinite loop. Perhaps humans could be thought to have a "timeout" function that quits the problem after so much effort is expended on it, but that is not a means for determining decidability.

We have just demonstrated that the set of strict, absolute, in principle unknowables contains at least one element, the undecidability of cognitive closure. Because cognitive closure is undecidable however, it could still be true that we have cognitive closure. It does not rule out that option, it just shows we can't prove it. *Which means that this set contains indefinitely many elements*. Because cognitive closure if true would mean that there is an unknown of quantity of things that might be over our heads. If the only thing that limits our mind is the realization that we are unable to know what is humanly unthinkable, then we aren't actually limited, we merely have a trivial blindspot so to speak but can see everything else. The set of unknowables would also have to contain

the proposition that it is undecidable that it contains any determinate number of elements, finite or infinite, because the answer to this question would presume that we were able to tell what lies beyond our cognitive closure. The cardinality of the set of strict unknowables in other words is itself unknowable.

The proposition that the set of in-principle unknowables contains indefinitely many elements, therefore, must itself also be considered undecidable and strictly unknowable if cognitive closure is undecidable. We also know that this set must also contain the set of undecidable mathematical propositions according to Gödel's theorems, which itself may be indefinitely large. If cognitive closure is true there could be an infinite number of possible truths that cannot be known because our minds can't know them, but the one we can count is the fact of cognitive closure itself.

Put another way, we are not able to know or enumerate the set of things we are unable to know. It's only possible to know we are unable to do so. We are able to decide this proposition about the innumerability of unknowns in the affirmative.

The language here admittedly gets a bit twisty. A decidably undecidable proposition is equivalent to a *knowably unknowable unknown.* An example of this would be, "It is unknowable what is outside the light cone at position x." Or "the living don't know what it is like to be dead." Or "the information state of a black hole's hidden region cannot be reconstructed." These are certain end points that line the boundaries of conceivability and which define, by the barricade they draw around it, the limit of the knowable.

That which remains absolutely unknowable under any changing circumstance is important to identify because all knowing occurs in a dynamical setting in which conditions shift under the input of new learning. What was relatively unknowable at time t becomes knowable at time t+1. Precisely what we will never see, never think about, never come across or more pointedly *can't* think about, in a way gives us what we do see, dwell on, and encounter. The entire edifice of number theory would collapse if the concept of zero was absent to provide the basis for additive identity. Among numbers it is atypical in so many ways, and despite representing the numeric equivalent of *nothing* it is a

definite something in terms of its considerable mathematical meaning. In the same way, cognitive closure, the limit, is the zero of our experience. The exception defines the rule.

As this chapter has shown there are certain conclusions we are driven by the effort to logically determine the limits of computability, and by extension, the limits of the knowable. The outcome of this pursuit is that this is an answerable question that we can know with certainty. This is not the luminous, all-seeing certainty of absolute knowledge. It is like being able to spot a black hole by seeing where everything outside of it ends, or to detect the pull dark matter on regular matter, or to notice the blind spot in one's visual field. We can identify the unknown by its contrast with the known but we do not thereby know it.

While the incompleteness theorems are profound statements about the limitations of logico-mathematical systems, there is no good reason to hold that we are either free of or subject to the same restrictions. If we are one way or the other, this answer is confined to the undecidability of cognitive closure. If we aren't, then as per Deutsch, we have computational universality which must be implementable in machines. To confirm it either way we would have to pit on an "alien thinking cap" and take off our human thinking cap to see the world with different eyes and think with different thoughts, to realize that there are some things that are just over human beings mental thresholds.

As has been demonstrated hopefully to satisfaction, the undecidability of cognitive closure, is provable, but only if we take it to be self-referential. Our own attempts to confirm it supply the proof that we can't. To be able to determine if cognitive closure is true, we would have to be able to experience a thought that is outside what is experienceable for human beings. Given that that is impossible, we can know that cognitive closure is itself closed. There is no way to demonstrate it, which self-referentially and paradoxically, is itself a demonstration of its proof. There is certainly a whiff of Gödel to this conclusion, Does it contradict or agree with his own belief in the capacity of the human mind to infinitely surpass machines?

Here is where things come to a head. Turing sought to develop an algorithm for a machine that could determine that any statement is undecidable (that could

decide it was undecidable). In congruence with Gödel, he ended up proving that such a machine was impossible. Yet we humans, as has just been shown, seem to be exceptions to this rule, meaning, perhaps as Gödel believed, there is some transcendental non-algorithmic attribute to our thinking. On the contrary it merely shows that, as Minsky pointed out, we exhibit a reflective property that allows us to detect our own states "from a remove", that is, to introspect on our own linear chain of reasoning from outside it. To use, in other words, metacognition to introspect on our mental states. Perhaps this has something to do with the parallel processing of the brain, where we have one stream of processes and another stream of processes that monitors the former. We can tell that our own knowledge is not both consistent and complete, but that is because we are under no obligation to be consistent and complete.

The gist of Gödel's argument is that we can look at a statement that cannot be determined to be decidable from within the formalism that generated it and realize it is undecidable. But there is no magic there. We are simply not confined to the same rules as the formalism. It's inability to determine that proposition is its own problem, not ours. Indeed, the same proposition rendered in another formalism could be decidable.

The following section is a technical demonstration of the above. It will provide a formalistic rehash of the argument for the sake of thoroughness.

(For both conventional deep learning models and newborn babies, the starting conditions are sparse, containing only the "genetic" framework that implements the possibility for learning. The learnability of the system is what makes it.)

Decidability is important for establishing knowability criteria because every unit of knowledge is a decidable proposition, that is it is knowable that is known. There is some way to formulate every item of knowledge such that it can be answered by a yes or know question. The limits of knowledge are therefore intimately wound up with decidability.

Let D = the set of decidably undecidable propositions; that is, any proposition

for which it can be decided there is a Turing machine which cannot prove it is decidable

Let G = the set of mathematical propositions applicable to Godel's theorems e.g that are true but not provable

Let C = the proposition that cognitive closure is undecidable is true

If C, therefore C\`: where C\` is taken to mean, cognitive closure is true but not provable because the question of cognitive closure is cognitively closed (undecidable).

Let K = the set of all that can be known

And let U = the set of unknowable unknowns, that is, all that cannot be known in principle. U consists of n -> ∞ elements. That is, indefinitely many elements approaching infinity.

C\`, being decidably unknowable, is an element of U.

Let's assume that the universe of discourse is defined as any descriptive proposition of fact and the set of propositions about those propositions.

In this case, D is the set of sets that consists of \{G, C, U } many elements, where truth function for each element of D, T(x ∈ D) evaluates to true where T(x ∈ D) recursively applies the operator "it is true that it is undecidable that x" to every element of every set of D. The set of all knowables K consists of the statement T(D), "it is known that every statement and/or every statement of every subset in D is undecidable." K also contains the complement of D: every proposition that is in principle decidable. In other words, K contains every in principle answerable question. But it does not contain, crucially, any element of D except for its intersection with G. Another way of saying this is that K means, "Every decidable proposition is knowable and the set of every decidably undecidable proposition is knowable. But no undecidable proposition that is not a Gödelian statement is knowable."

The set of all possible knowables is the set of what is known to at least one knower or which is not an element of D excluding its intersections with G. While the truth function of D itself, T(D) is taken to mean "it is true that every element of every element of D is undecidable", is an element of K.

The complement of K, clearly, is U. This does not pose a contradiction even though T(D) is an element of K. Because the set itself D is knowable, while its subsets and subelements may not be. Another way of looking at this is that U contains every question that cannot be formulated and every unanswerable question.

Suppose that K is finite. The size of K is persistently smaller than U, since U is a member of D and approaches infinity (cannot be enumerated), even if most of the elements of U are trivial. U cannot by definition be enumerated because there is no way to precisely quantify the scope of unknowns. The number of propositions in U is necessarily of indeterminate size but may approach infinity. Presumably if the universe is finite K must also be finite. Supposing the universe is finite, all empirical knowables regarding facts about the universe are also finite, allowing that we exclude uninteresting regress statements such as "…it is true that x and it is true that not x is false and it is true that this whole statement is true" and so on.

But U also contains a set of counterfactuals which is also innumerable. We cannot know what might have happened had some historical event never happened. We cannot know what would have happened if Julius Caesar was never assassinated or WWI never transpired. Presumably, however, there is a factual answer to these questions in principle. But we cannot know what might have happened in those possible worlds where events would continue to branch off in untraceable ways, compounding the counterfactuals many levels deep. There is no experience available to us that would give that knowledge, even though, presumably, such an experience could have become available had things happened differently and we inhabited that timeline.

Allowing that counterfactuals are unknowable whilst also being valid epistemic entities greatly expands the size of U and makes it larger than K.

Therefore, more is unknowable than what is knowable and some propositions that cannot be known can be accorded that status with certitude.

Here we have gone to some effort to show there is a discernible set-theoretic structure to all epistemic contents. By adding a modal disposition to epistemic contents, an *-able* predicate, we can establish the order of these sets. The set of

unknowables contains the proposition that there are indefinitely many unknowables, with this own statement being a member of the self-same set. Yet the cardinality of the complement of this set, all knowables, is itself unknowable and belongs to U, because its size is logically bound to the set of unknowables (it is what left when all unknowables are removed from the superset of all possible epistemic considerations). This cannot be determined however because the set of unknowables cannot be assigned a cardinality. Presumably, the uncertainty principle in physics applies to every atom in the universe, which means the proposition, "the position and momentum of this particle x are unknowable at the same time" appears at least 10^79 times in U, for each atom that is estimated to be in the universe. So the set of unknowables is quite large just going off this as a base rate, even if it is in principle uncountable.

While set U is vast, it contains many trivial elements, perhaps infinitely many. K, though while smaller than U but still very large, is more compact, although it also contains many trivial elements. It contains statements of natural law, for example, that specify regularities that apply universally. There are certain "golden statements" within K that give you more bang for you buck, by specifying a law that captures the causal essence of a vast array of phenomena. Even the statement, "the position and momentum of this particle x is unknowable at the same time", when generalized to, "the position and momentum of every particle cannot be known at the same time" is knowable, and serves as a benchmark for these kinds of compact universal statements. Thus, K contains universally quantified statements which describe natural laws.

Again, we know that U is larger than K because U contains infinitely many or at least far more counterfactual statements that are unanswerable relative to our timeline. Our timeline is merely a subset of these eventualities, one branch off the tree. Some may deny that these counterfactuals have valid epistemic status, that for example, the claim that history would have unfolded differently had Caesar never been assassinated is not unknowable but meaningless. The word itself "counterfactual" means against the factual, and we are committed to defining the universe of discourse as the realm of the factual, that is, the actual rather than the merely potential. This supposes however that the causal structure of the world is

not rigorous enough permit would-be facts had things turned out differently. There may indeed be a parallel universe where these things did happen. We can't know that there is and and we can't know there isn't.

I am inclined to believe that there are such things as factual counterfactuals. If Frank did not drink the night before, he would not have been hunger over. If John Wilkes Booth had gotten a fever, or had a change of heart, he would not have shot Abe Lincoln. If so and so did not smoke a pack of cigarettes a day, they would not have gotten lung cancer. These sorts of claims have more than a mere air of speculative suspicion about them: I know them as well as I can claim to know anything that is not self-evident. I know them to be true as much as I am willing to believe in causality.

While these claims would therefore seem to belong to K, what does not belong to it is the world where John was not hungover, or Abe Lincoln stayed president, or this person abstained from cigarettes. Hence, if we can accept that anything is true about the merely possible, we must commit to the truth of these claims.

Certainly history would have been different in the minimal sense that Caesar would have continued to live for a time, but that is all that can be known for certain about that timeline. But for every statement of fact whose negation is not bound by logical impossibility, we would have to ask ourselves the question if it could have happened differently and if so, what might have happened. Each and every such question is undecidable. Every such alternative timeline specifies its own uncertain future, each of which is mostly unknown, and thus vastly amplifies the space of unknowns.

Part III: The Modal Logic of Ignorance

Beyond the physical and computational layers circumscribing ignorance, there is also a deep modal level. What is possible to know, and moreover, how does possibility relate to knowledge and ignorance?

Given that the uncertainty of an event is partly function of its novelty, it is helpful from a conceptual standpoint to do away with those ontological

restrictions which force us to regard an event as purely novel by situating it in the fullness of possibility. If we embrace a rich possible worlds semantics it may allow us define the full scope in which uncertainty in its true measure is really not uncertain. Already modal reasoning is industry standard for epistemic logic, and if anything ignorance is more closely associated with possibility than knowledge. To know is to constrain possibilities. Ignorance, as the correlate of maximum entropy, is also the maximum of degrees of freedom in an information system's phase space, the maximum of possibilities. For any ambiguous proposition there corresponds a possible world in which each of its interpretations has satisfiable truth conditions. In which case such a proposition is knowable, else it is strictly unknowable. If every possible event, no matter how ambiguous we are about the probability of its occurrence, has at least one possible world in which it is actual, suddenly we are able to dispense with an impoverished conception of our world as the one and only standard of comparison. The relations between ignorance and possibility become manifest. In this respect, the controversial but energizing work of David Lewis supplies fruitful soil in which to grow the concept of ignorance.

The Many Worlds of the Unknown

What is the connection between possibility and knowledge? To know is in some sense to declare an invariance, to assert necessity or at least to privilege a subspace of all possibilities as more probable than the rest. If you know that it will happen, it necessarily must happen or else you didn't know. Ignorance, as a function of entropy, signifies a lack of stipulations about what is possible. Possible, to be clear, not in principle but in practice, because the strongest modal criteria include what is possible in all or some possible worlds. Pragmatically we are really concerned about what is possible in anything but the actual world and it is only when examining the deep logical structure of epistemological fundamentals that we are concerned with possible worlds. To be ignorant is to be thrown open to arbitrariness, to lack any means to rule anything out.

The catch here is that there is little means to specify as a general rule what is

practically unknowable or knowable, given the great degree of variance that is presupposed by an arbitrarily defined knower. No universal, context-independent knower can be defined by which all cases of practical knowability are articulated, because any finite knower has their own "epistemic profile" which enables or disables them from knowing things relative to their given state and spatiotemporal relation to the fact. So there are no strict rules for practical unknowability, only conditions for information access that admit some possible knowers and exclude others. Unsolved mysteries do not imply unsolvable mysteries. They are only unsolved because the conditions for knowing them are unmet. Unless the laws of nature were defied in the course of that mystery's miraculous origins, then it is in principle knowable. In terms of precedence, everything that is in principle knowable is knowable in practice but not vice versa. So by prioritizing in principle knowability, we also characterize the space of pragmatic knowledge. Consequently an adequate understanding of possibility is a prerequisite for an adequate understanding of ignorance.

An ambiguous proposition can be thought of as analogous to a semantically "high entropy" statement. The more ambiguous a statement generally the more uncertain it is, the more possibilities, that is, fit its semantic scope or logical form. In this chapter we are concerned with how ambiguity and possibility relate to the definition of ignorance we've been developing.

To be aware you don't know something is often met with two responses. You either directly seek information about it if you have reason to believe answers are available, or you speculate about what could be the case based on your present background knowledge. If you chose the latter option, you begin to enter a state of mind we call wondering. You cycle through various hypotheses and reason about the merits and flaws of each of them. To cycle through various hypotheses is to picture various conceivable possibilities. Ambiguity, therefore, is an invitation to engage in modal reasoning: to contrive what-if scenarios in search of the most plausible account for what's under consideration. Possibility leads to plausibility which in turn leads to probability, as we hone our conception of the truth down from the general to the specific. A hypothesis must first pass general

laws which define what is possible to be considered as a candidate explanation. It must be possible. It must further pass various evidence criteria and restrictions appropriate to the specific context of the circumstances in play. It must be plausible. Finally it must have a certain likelihood of being the case given the various conditions and frequencies of such events as they were known to occur in the past and given what the inclusion of prior evidence says about what posterior probabilities ought to be expected. It must be probable. Possibility here is simultaneously the strictest and loosest criterion, as what is possible given certain laws is a vast set of combinations, but also, what is impossible relative to those general rules must similarly be vast. Implausible and improbable events do indeed occur, it is as Hume said "by custom alone we think the future conformable to the past." Nevertheless they are good heuristics by which to help fix an undetermined propositional truth value.

The key here is that you know you don't know. You're aware of a gap in your knowledge which means that gap is contextualized by everything you do know. So even though you are ignorant of that missing piece, you can conjecturally project the outer context of knowledge onto the gap. There are so to speak "nerve endings" terminating at the edge between your knowledge and the ignorance gap, suggesting there is some definite formation that can couple with those "afferent" connections. You might wonder if Shakespeare really wrote the plays attributed to him, but you know for a fact it wasn't Keats who did it; he wasn't born yet.

There is then a very peculiar correspondence between the structure of possibility and the structure of ignorance. We seek to locate the one truth in an array of plausible counterfactuals. An event that is both possible and plausible is thus a candidate for being true.

Lewis' work is distinctly "far out there" in terms of the lengths it goes in ramifying the idea of possible worlds. Of particular interest is his work on nomological necessity ("Counterfactual Dependence and Time's Arrow" (1979c)). Effectively what is possible in a given world is restricted by natural law, and there are families of possible worlds that conform to those particular laws. As such, the region of plausibility in our universe is circumscribed by how the laws

dictate the combinatorial possibilities of our world and those within the family of worlds to which the laws apply. To the extent that we have knowledge of natural laws, we draw a line between the impossible and the possible. This is a crucial piece of knowledge, it is the first approximation of what is true. Natural law knowledge of a special class, a profound form of knowledge that tells us what is possible to know about specific causal histories in the universe to which unfold from it . Natural laws are means by which we can define the rules which specify every possible configuration of everything in a world. Whatever happens must happen according to natural law, supposing we have the correct understanding of them. According to what Lewis called the principle of recombination, natural law circumscribes the domain in which discrete events can coincide with others. "anything can coexist with anything else, at least provided they occupy distinct spatiotemporal positions. Likewise, anything can fail to coexist with anything else" (Lewis, 1986, p 88). A world is defined by the set of possible recombinations of its entities. Insofar as a given event transpires in such a world, it must belong to a subset of those possibilities. In which case, we can be said to be meaningfully ignorant of a proposition if its semantic payload designates a true possibility belonging to the world that we are unaware of. If it suggests something impossible then not knowing it carries no significance. Given such and such conditions, let's say, that rule out the possibility of one falling through a crack in the earth all the way to the earth's core or breaking free of earth's gravity and floating away, our ignorance of what this eventuality entails is of no import. Caution about the possibility of slipping on ice that formed following a warmer day that melted snow and then refroze that night, however, is a live possibility and thus a going epistemic concern. Ignorance here could have serious consequences. The possibility of an outcome, and thus by degrees its probability of occurrence given those conditions, assign a range of plausibility values to unknowns, which in turn specify the degree of importance of those facts.

This is important for several reasons, but the most relevant is that what contradicts natural law is not something we should concern ourselves with. There could be an infinite list of statements which contradict the natural laws of our universe, and not one of them is worth the bother to know unless we are

distinctly preoccupied with understanding the combinatorial *possibilities and necessities* of different natural laws themselves. We can imagine a universe where the value of the gravitational constant was different, and this may be a useful exercise for the physicist. A computer simulation of a toy universe might wish to experiment with altering the constants that govern its model of physics. However, so long as the constants are what they are for our world, it allows us to draw a definite line. Subtracting an infinite number of items from our laundry list of epistemic concerns is a good bargain. In the immortal words of Sherlock Holmes, "when you've eliminated the impossible, whatever remains, however improbable, must be the truth." Of course in this situation one must have certain knowledge of the impossible to arrive at the possible as the remainder of that exclusion, but we assume that in some circumstances it is possible to know what is impossible to a high degree of probability based on the given evidence and inference from the best explanation.

Before proceeding further the definite relation between possibility and ignorance must be more firmly established. The link between them is ambiguity. When we are uncertain of an outcome it means one of two things. Either we have precisely no idea at all what will happen; the outcome and its cause is beyond reckoning. Or we are cognizant of a number of possible outcomes but unable to decide which of them will come to pass. The act of considering possible outcomes is itself an expression of ignorance as to what will definitely happen, given that certain knowledge eliminates any option of alternative possibilities. Whether a given event will occur in reality to comes down to a definite yes or no, rather than a decimal fraction up to or between 0 or 1 as used in calculations of probability. Put another way when one is certain and correct there is no possibility that the outcome could vary, the outcome is *necessary*. Though we can indeed be uncertain of a necessary cause as we might a contingent one that affords several possible outcomes, insofar as a necessary cause derives from a natural law, as Lewis points out, we do not have to concern ourselves with tracing the particular causal history of the event but can refer it to the law that governs it. Contingent events, however, are more bound to their distinct causal histories. Otherwise, we can fall back on the general principle. Contingent events

have no such "gold standard" backing in natural law, at least not in any way that can be immediately recognizable, and so it is sensible to assume that the outcome could be one of several seemingly arbitrary possibilities. In other words, natural law does not forbid multiple variable outcomes of the given event but it does place an upper limit on that range.

Ambiguity is precisely nothing more than a state of semantic indetermination in which the meaning of a proposition hangs in the balance because it lacks any single necessary interpretation. The question of import is then: is the proposition intrinsically ambiguous or is it our propositional attitude towards it that causes it? "It may rain Thursday", expresses a measure of indeterminacy, but epistemologically it can be translated as "I do not know that it will in fact rain or not rain Thursday." Thus, we signal our ignorance by using a modal verb that bundles together several indeterminate possible outcomes. Are there propositions which necessarily scope to several possible truth conditions, is possibility ontologically "real" in some cases? Or is it only our subjective indeterminacy that cloaks necessity in an illusory garb of contingency through modal language? The immutable fact is it either will or will not rain Thursday, there is no maybe with any ontological credence. This is no small distinction. If this weren't the case, and let's say the world "made up its mind" each day to decide whether to rain or not, it would mean that there are *external causes* of ignorance, that our ignorance-as-ambiguity is not merely a deficiency on our part but that it is induced by the loose and arbitrary combinatory powers of the world. There is, then, inherently ambiguous events and modal language simply taps into it, *denoting* it. If the later is true, modal language merely expresses subjective ignorance about necessary events about which we have no knowledge of the definite causes.

Modal propositions revealingly seem to express a weaker grade of knowledge by watering its meaning down with uncertainty. "It may rain Thursday" seems to play it safe compared to "It will rain Thursday" assuming it does in fact rain. A modal statement covers its bases, hedging its bets against the possibility of being wrong, but in doing so it dilutes its meaning. "Anything could happen" while necessarily true, is vacuous. To the degree that a statement makes a

concrete assertion of a definite outcome, "*This* will happen", to the extent that we can verify its truth conditions by specific criteria, that statement will come across as more knowledgable than the ambiguous one. For the clear reason that the ambiguous statement bundles *more ignorance* into its assertion by committing itself to an arbitrary array of possibilities rather than to a single confirmable assertion. "Anything could happen," makes no pretense of knowledge, though it is vacuously true; while "This will happen" connotes certainty and, if not knowledge explicitly, because after all it may or may not happen for reasons unrelated to those supposed, it does express less ignorance than the modal statement. After all, one can simply guess without knowing that it will rain Thursday. However the measure we are trying to take is the degree of ignorance the proposition expresses, not the amount of knowledge. If false, or even if true for what was believed for the wrong reasons, "This will happen" can be said to be wrong by only one degree of ignorance. "Anything could happen" while necessarily true, regardless of its outcome expresses the highest degree of uncertainty, a kind of semantic maximum entropy. Its lack of commitment is an admission of ignorance, equal in amount as to how many eventualities are designated by "anything." A statement which is always true and has no possibility of being false is an empty truism. A statement which has definite truth conditions has "stakes" on it, it does not assert an endless variety of possible outcomes but confidently seeks to pinpoint one and only one. Such a statement is more airtight and has less ignorance implied by it than an open ended modal statement.

One might claim that had Archduke Franz Ferdinand not been assassinated WWI would not have happened. Can there be any truth—or even falsity—to this claim, or is it strict nonsense? Our instinct seems to tell us that there is a meaning to it, that such counterfactuals are at least sometimes worth entertaining if for no other reason than for the intellectual exercise of considering how history might have played out differently. Perhaps had he not been assassinated, the gears rolling toward WWI would not have started turning on that fateful day in particular but would have started up another day sooner or later. Perhaps the manifest potentiality of WWI was latent in the sociopolitical tensions of Europe

at the time and some other precipitating focal event would have inevitably lit the powder keg. Perhaps it is indeed true, and that the whole travesty of the war could have been avoided had the bullet that killed the Archduke never left the barrel of the gun that would've shot him. Had he not been assassinated and WWI would therefore never happen, might'n't WW2 also have never happened? And so on, one can continue traversing a chain of counterfactual dependencies that seem to hover above our actual world, tangental and vaguely informative about it, yet unreal, purely speculative.

Yet at the same time our gut tells us that there is some substance to these counterfactual statements, even if they are unanswerable. The key point here is not that these questions are meaningless, but that we are ignorant of their answers given the hard fact that our timeline excludes confirmation of those possibilities. If we could replay the tape of history like we could replicate a laboratory experiment, ensuring this time the Archduke survived, we might just as easily confirm our counterfactual hypothesis as if we were carrying out a routine scientific trial. These are questions that could in principle be answered if there is in fact an alternative world-branch in which such eventualities unfolded that could be tested in our world. While granted this modal interpretation of ambiguity has its critics, (Devitt 1976 gives a particularly scathing take down [https://www.jstor.org/stable/27902433?seq=1]. See also Kuczynski, 2007 [https://www.sciencedirect.com/science/article/abs/pii/S0378216606001822]) I find it to be an attractive account given how it mirrors the notion of entropy in semantics.

To the degree that we afford possibilities first-class epistemic status, as argued by Lewis, we dispel ambiguity, or rather map out its alternative interpretations each of which is necessarily true for a given world. Making the task of interpreting an ambiguous proposition's meaning for us an issue of selection, choosing the one alternative the fits our world. All truth conditions of the ambitious proposition, that is to say, correspond to a possible world, and thus the effort becomes to discover which one belongs to ours. In the fullness of possibilia there is no ambiguity, only valid counterfactual histories that terminate in said statement shaking out as true in one world and false in another. There is no

exclusive OR, only a continues conjunction AND which sorts out each possible truth condition for each possible world, all of them being true depending on which world you quantify over. Of those many alternatives, all but those which belongs to our world's causal history would be considered false knowledge if affirmed about our world, though in the appropriate world-context, the opposite would be true of it. Only that alternative is answerable, which fits into our world, must then be true, and about which we can claim knowledge by acquaintance.

This raises an important point. Lewis questions if it is strictly necessary that we can only know by causal acquaintance. Ever since Russell refined and popularized the distinction between knowledge by description versus knowledge by acquaintance in his influential essay *On Denoting*, it's been generally accepted by the philosophical community that one of the valid ways of knowledge is by having sensory familiarity with the object. To know something on this account is to have been caused to know it by experiencing the object. Lewis singles out the criticism that the truth conditions of possible world semantics are unverifiable, because we can only inspect our world, not those of merely notional possibility. He counters these objections by comparing modal knowledge to mathematics. (Lewis, 1986). Mathematics entails an infinite number of truths with which we are not causally aquatinted, for instance probably no one has taken pi to the trillionth digit, divided it by six and added one to it. Mathematics gives us knowledge of arbitrarily many mathematical objects which have not been directly subjected to inspection. Certain esoteric areas of math, for instance non-euclidian geometry, describes shapes which are not approximated by anything made available by acquaintance to us in our visual field.n Lewis wonders why, if this is held to be true of a mathematical language, why would we not hold it true for non-mathematical language? Modal reasoning, held to the same standard as mathematical reasoning, can therefore speak to objects which exist within the context of possible worlds.

Why is this way of looking at things insightful? Because it helps us to situate knowledge and ignorance in a logical plenum, a manifold in which everything is accorded its proper place, whether true or false of our world, whether correctly or incorrectly believed. To the degree that vagueness or uncertainty clouds a

proposition, there is a semantic entropy which inhibits us from discerning which possibility out of several applies to our world in this instance. I suspect that the reader who is left cold by Lewis's modal realism will not be swayed by my stabs at connecting it to the theory of ignorance. Be that as it may I stick to the assertion that ignorance and possibility are irreducibly linked. Having scouted the relationship between ignorance and possibility from a logical standpoint, the next stage is to determine how the possibility of knowledge is realized by acts of discovery.

We are about to enter quite a complicated subject. The reader should be warned that this section will be quite hard to follow. It demands as part of its very subject matter labyrinthine chains of the words known, unknown, knowable and unknowable in various combinations. There's no way around this, so I will endeavor to make it as clear as possible given these difficulties.

The following will make several assumptions. First, that what is knowable/unknowable is amenable to a set-theoretical analysis. For this to be the case, what is knowable etc must be discretizable—it must be capable of being reduced in terms of finite logical elements. Now for this to be the case, it is not enough to assume that the facts themselves are finite, rather, we can with certainty claim that any proposition of fact must be finite in the sense that it takes less than an infinite amount of time to evaluate it. Therefore, we will speak of sets of knowable and unknowable factive propositions. The facts themselves may be infinite in detail or concern facts about infinities, but every knowable or unknowable proposition of a finite or infinite fact is itself finite.

A specialized language must be employed to analyze the possibility of knowing. Just as there are neg-epistemic ordered pairs, there are also modal neg-epistemic combinations : knowable knowns, knowably knowable knowns, knowable unknowns, knowably unknowable unknowns; unknowable unknowns, unknowably unknowable unknowns , unknowably knowable knowns. These tuples, or ordered collections of a given size, are difficult to parse at first glance, but with some explanation and examples to go with them become much clearer and their logical utility more apparent.

Knowable knowns and knowably knowable knowns risk coming across as uninformative tautologies. They are indeed but they are necessary to close the loop. All knowns are a priori knowable. Knowable knowns are counted here strictly as a formality. However, with regards to speaking of what is knowably knowable, it is meaningful to recognize that a knowable fact can be known, otherwise one might overlook or discount it as a dead opportunity for knowledge. It is possible to speak of an unknown knowable, or surprise, accidental knowledge.

These modal combinations express up to three layers of meaning. First, whether a given fact is known or unknown. Secondly and thirdly, whether it is knowable/unknowable if said fact is knowable/unknowable. These are distinct propositions.

Modal tuples are necessarily conjectural in a way that basic pairs aren't because they affirm categorical rather than contingent relations. A unknown unknown is different from an unknowable unknown in that it has the possibility of becoming known. Rather, unknowable unknowns are a subset of unknown unknowns. Establishing grounds that it is impossible to know something is harder than establishing that one doesn't know. Even so, the modal forms are also more powerful in this respect. If true they assert universal facts about the conditions for knowledge.

An unknowable unknown differs from an unknowably unknowable unknown by a degrees of certainty. The former is certainly unknowable meanwhile the latter spells a lesser degree of certainty because it can't even establish that a fact can't be known. That something is an unknowably unknowable unknown means that one can't know you can't know is is unknown. It does not forbid the possibility of knowledge but makes it doubtful.

These modal combos are useful when reasoning about, in Kant's phrase, the conditions of possibility for knowledge and for analyzing the question of strict unknowability.

As Rescher (2009) mentions, so called agnoseology, the study of unknowability, is fraught with logical paradoxes. There are various classification issues associated with how to sort what goes where when one knows something

can't be known. True unknowability is absolutely in principle unknown, such that no possible knower in any circumstance can know it in any possible world. How can you identify such an unknowable given your lack of familiarity with the premises involved and a lack of omniscience? If you can't know how you know it is unknowable, then nothing is proven.

To make matters worse, suppose one had two bins, the knowable and the unknowable. Where should one place a purported fact that is a known unknowable? Should it not belong to both sets? But that's a contradiction. Rather the proposition must be split into two distinct atoms, namely "it is known that p is unknowable" and "p is unknowable." To be unknowable does not mean to beyond awareness, otherwise the only unknowable is what is unknowable, it would be an unknowable unknown. Whatever an unknowable fact is, it must be at minimum describable otherwise there would be no referents to place into the unknowable bin. The set of all unknowables would otherwise be undefinable. Thus all truths would be knowable because nothing can be denied to be knowable, implying a strange sort of omniscience we aren't supposed to have. Another way of putting this is to say it is impossible to place anything in the category of the unknowable without making it known. This gets to the root of Fitch's notorious paradox of knowability. Fitch demonstrated that by applying the inference rules of epistemic logic, if all truths are knowable, then all truths are known (Fitch, F., 1963. "A Logical Analysis of Some Value Concepts," The Journal of Symbolic Logic, 28: 135–142; reprinted in Salerno (ed.) 2009, 21–28). This conclusion therefore either results from unsound premises or is false because not all truths are known, therefore not all truths must be knowable. But the reversal of the argument is clearly an inductive fallacy.

As shown in chapter 6, these paradoxes can be resolved or at least greatly clarified by distinguishing between assertions of knowability with second order reflexive assertions about these first order assertions. This permits disambiguating between where to place known unknowables… in the mutually exclusive categories of unknowable or knowable etc. To state something is a known unknown for example is really to state a metalogical compound proposition consisting of two separate conjunctive assertions, with the outer one

being a truth claim about the inner one. In this case the compound proposition can be sorted into the knowable set. The inner proposition taken in isolation, the unknown, however, cannot be treated the same and sorted into the unknowable. This lack of symmetric parity of inferential salience recurs frequently between knowledge and ignorance. What is known is a priori knowable, but what is unknown is not by that fact alone unknowable.

This incongruence hints at an important distinction between the sets of the knowable and the unknown. The set of knowables has the property where it intersects with those unknowns which are knowable, thus it subsumes all knowns and some unknowns. Whereas the known is a subset of the knowable, the unknowable is a subset of the unknown. Precisely this inversion of subsumptive relations explains why the unknowable is difficult to define. As the relative complement of the knowable, it is completely disjunctive from it. Whereas the knowable is the absolute complement of the unknowable, overlapping with a subset of the unknown. The knowable therefore subsumes the unknown in a way that the unknowable does not subsume the known. The unknowable is completely aloof form this conjoint superset.

Issues of self-referential unknowability should also be considered. Assertions about the size of the set of all unknowables are itself unknowable if the size is unknowable. Thus assertions about what is unknowable build out the set. Meanwhile the set of knowables is much more easily constructed, and is even extended as a byproduct automatically by searching for what is unknowable. For example as a freebie we know that the size of the set of unknowables is itself unknowable, because it is not self-indexing. The set of the knowable is indexed by what is known, but this does not happen in for the unknowable. Because if we knew what the size of it was, we'd have enumerate every unknowable, but what is unknown does not index the unknowable because they they can be disjunctive. For every unknown we might hope to identify, we also add a compound proposition in the set of knowables saying that it is known that it is unknown. Interestingly this has the implication that the set of the knowable *must be larger* than the set of the unknowable, because it includes all known facts, all knowable unknowns, and also all assertions about what is known to be

unknowable.

It would seem then, that some of the most confidently assertable unknowable facts are those propositions which themselves reference the scope of the unknowable. These sort of metalogical propositions are a bit of a sleight of hand, but they capture some fundamental details about how the unknowable is logically structured.

What must be true of the member must be true of the class. If the size of the set of all unknowables is unknowable, it is also unknowable whether the size of that set is 300 or 300,000,000 or any finite number. This however creates a seeming paradox. For every expression of the unknowability of the size of the set of unknowables, we add another proposition to the set, and this could be repeated infinitely. "It is is unknowable that the size is 2030, 2031, 2032… $n \to \infty$" Which would therefore imply that the set of unknowables is infinite simply on the basis that not knowing its size also means not knowing any of its possible sizes. If the set is contains an assertion for every finite number however, the set is therefore countably infinite. If it is infinite, it is then known that all of the previous propositions about its finite size are false because it is known that the size is not finite. Yet this happens to destroy the very basis for its infinity. It cannot contain an infinite sequence of propositions each asserting it is of an unknowable finite size without negating them all once having become infinite. But if it negates them all, it must then be finite. However not a single finite number can be known for its size, hence the paradox of how to stipulate "how many" unknowable facts there are based solely on the fact that how many there are is itself unknowable for each number.

A similar paradox arises when we consider whether it is knowable that the set of unknowables is empty . It may be the case that every fact can be known, but it's not clear how we can know that. It is a known unknown. If the fact itself can't be known, an unknowable unknown, it is false. If it is true that every fact is knowable in principle then obviously the set of unknowables is empty. But this is the only condition in which this is true that the set can be empty, and it is not clear how to demonstrate the conditions that make this true.

The problem of cognitive closure entails that the limit of knowledge or

thought cannot itself be known or thought, and this paradoxically embodies the limit it denies. However the fact that closure is not decidable does not mean such a limit exists. Although closure cannot be answered true or false, it is just a limiting case, and if it is false, it still does not strictly mean everything may be knowable. If it is true, it does not imply anything more than that the mind cannot think outside its capacity, but that does not mean that anything unknown falls outside that range and is unknowable. There can be no knowledge without a knower. What is unthinkable does not correspond exactly to what is unknowable: a squared circle is unthinkable, neither is a natural number smaller than zero, but such things are not possible to know because they are not facts.

Unless it is true and known that everything is knowable it is unknowable that the set of unknowables is empty. The set of the unknowable would then contain a statement saying that it is unknowable that the set is empty, which is self-contradictory. It may very well be empty, but it cannot be empty and unknowable that it is empty at the same time. It cannot be the case that it is unknowable that everything is knowable without contradicting complete knowability but it can be the case that it is unknown that everything is knowable. The knowledge that everything is knowable cannot be demonstrated even though it may be true. It cannot be demonstrated for the extralogical reason that now matter how much is known it is impossible to calculate the remainder of the unknown. This option leaves open the possibility that the the set of unknowables is unknowably empty.

Thus in these situations the set of unknowables can neither be infinite nor empty without contradiction. Usefully from a logical perspective this helps us to characterize its extrema, while this does not prove it is not infinite or non-empty, it at least lends a logical order it.

But this jumps too far ahead. We should concern ourselves with the minimum condition, the identification of at least 1 unknowable proposition, rather than wading too deeply into the thickets of infinity, circularity, self-contradiction, and self-reference regarding the unknowable. Although it should be stated that these kinds of self-referential concepts are common in set theory and have even been applied to attempts to establish the foundations of mathematics. So there is

precedent for using them, and they are strictly logical criteria of unknowability that are applicable as a priori reasonings independently of other concerns. These merely logical statements are facts among others, just as unknowable or knowable as any material fact and no less real.

Though it is valid to include reflexive statements about what is unknowable about the unknowable in the unknowable, there is a certain tautological vacuousness to such a practice. This is not what we're really after. Putting it aside, these self-reflective logical propositions are of secondary importance compared to unknowable facts about the world. Part II section II went into detail about the question of cognitive closure and how it was undecidable. It is perhaps the keenest example of an unknowable unknown. This section concluded with a proof that the set of the unknowable has a least this one element. To know what we can't think is to think what we can't think. It is like stepping outside space. There is always more space ahead of us and it is presumed by the very act of stepping.

What is of interest is in principle unknowability. For any such proposition, we should be able to state why exactly it is unknowable in all possible circumstances for all possible knowers. It presupposes strict impossibility. The in practice unknowable is always conditioned by relativity, it is unknown with reference to a circumstance and a subject, rather than because of some absolute necessity. It is quite easy to establish a set of relatively unknowable facts, because for any given individual reasons can be specified why this fact is unknowable for them. I believe I have solid grounds for asserting that it is unknowable to me (at this moment in time, assuming I don't get abducted tomorrow) that there is other intelligent life in the universe. In principle however this must be knowable if barriers of space and time were removed. For absolute strict unknowability it seems impossible to specify both what is unknowable and the conditions that make it so.

Perhaps it is only feasible to specify logical impossibility as conditions for strict unknowability, for the reason that no one observer has access to all that is knowable and from that can compare their knowledge with the world to see if there is yet something unexplained. To claim that something is unknowably

presumes omniscience, otherwise there is no way to rule out that the unknowable is only so relative to an insufficient knowledge base or incomplete information access. One might hold that the principle of sufficient reason means that every phenomenon has a rational explanation, that every fact is explainable and therefore knowable. This is associated with what Chalmers (2012) (David J. Chalmers, Constructing the World, Oxford University Press, 2012, David J. Chalmers, Constructing the World, Oxford University Press, 2012) calls the scrutability thesis:

There is a small set of truths, with conjunction D, such that for all truths S, ⌜*If D then S*⌝ *is a priori.*

There are in other words, implicit connections between facts, and that from a set of fundamental facts one is able to derive all facts. In terms of what basic facts one would need to derive the total facts about epistemic possibility, the lack of knowledge of all physical facts and all facts about intelligence rules out what defining what is physically unknowable or unknowable to all possible minds in the physical world. Supposing these restrictions are true of every mind in the physical universe, however, ironically it would then seem that one unknowable fact is that physical and cognitive unknowable facts are unknowable.

For these reasons it can be maintained that epistemic impossibility is best characterized by purely logical relations and facts directly pertaining to what is knowable and unknowable. The conclusion is that there are indeed certain to be unknowable facts, but these facts are not empirically derivable facts about the world. Rather they are concerned mainly with the logical connections that hold between the knowable and the unknown and the unknowable in relation to itself. In the absence of perfect information however, which is likely true of all possible finite knowers, the possibility that there are unknowable physical facts cannot be ruled out. For example, the position and moment of a particle, the position and momentum of two particles or all particles, and so on.

A useful concept for picturing the dynamical relationships between knowledge, learning, and ignorance and the possibilities in which entwine them is what I call discovery space. While I have alluded to it elsewhere here I will

give it a solid definition.

Discovery space is an epistemic framework that encodes information about the proximity and relationships between certain knowledge states, learning steps, and states of ignorance. There are certain dependencies that must be satisfied for certain unknown facts to be made discoverable. Math is a concise demonstration of this fact: one cannot learn calculus without at first learning algebra and one cannot learn algebra without first learning arithmetic. Knowing certain things makes learning other things more likely. Certain "regions" of knowledge are contiguous with certain regions of ignorance. To get to calculus, one must first satisfy at least two prerequisite learning steps. Passing these steps of course consist of passing their own learning steps, all the way down to the most elementary facts. To learn arithmetic you must learn subtraction, addition, multiplication and division and so on. Learning certain facts often gives you another fact for free, this fact is adjacent to it in the discovery space. For example, addition and subtraction or division and multiplication are inverse operations respectively. Understand one and you are equipped to understand the other.

Here, the modality of knowledge, that is, between know-ability and unknow-ability, is a critical distinction. We have spoken of known knowns, known unknowns, etc, and this is a useful manner of speaking adopted widely in the literature on ignorance. But adding a modal modifier to these terms changes the dynamic substantially. To go from known unknown to knowable unknown and finally to unknowable unknown, describes progressively restrictive transitions of modal actions where valid "moves" can be made or where there are permanent blocks that limit the space of moves. Where one stands on the "board" of discovery space at a given state, determines what is knowable at that step.

To find out is to be positioned to find out. The discovered datum is near—if not literally near in terms of spatiotemporality than near in the logical sense of the current knowledge state being logically connected to that discoverable fact. Our own bodily experience prompts us naturally to think of discovery as searching through a space. One can picture the process of discovery schematically as a two-dimensional grid consisting of various cellular tiles. The agent, like a cellular automaton, is positioned in the middle of the grid. It has

limited "vision" into the adjacent nearby tiles above and below and to the right and left of it. At first let's say it starts in a state where the agent can only see the four adjacent tiles to the right and left and to the top and bottom of it. Beyond that is a "grey area" of unexplored tiles. Whether the grey tiles are knowable or unknowable is unknown, they exist in a kind of Schrödinger cat-like state of indecision. As the agent moves from tile to to tile, its vision expands to new tiles. It discovers that each new tile it encounters is either crossed off with an x, meaning that it is inaccessible, or not crossed off meaning available to move to, a free space. It can only move in the free spaces but cannot pass through the x's. The free spaces represent knowable unknowns, moves that the agent can make as it navigates the space. The crossed off tiles represent known unknowables. The agent can know it can't make that move or discover what occupying that tile reveals beyond it.

Let's assume as the automaton navigates the grid by making available moves, its vision increases so that the true state of more gray tiles are more and more revealed. Each tile it crosses becomes "colored in" as knowledge, allowing it to see the tiles adjacent to it do a limited degree of depth, such that the more tiles it traverses the wider its field of vision. However, for any spaces that are blocked off by a ring of crossed off tiles, everything inside that space is an unknowable unknown. There is no means to access it, despite the agent's explorations providing it with the ability to see the outlines and contours of where it begins and ends. These closed off spaces are sort of like a black hole's hidden region. We can see where it ends, but we can't see inside it.

This little game helps graphically represent the modal relationships between epistemic states using a spatial framing concept. If motivated to do so we could probably add a temporal dimension to it as well and with it a concept of causality. As the agent moves around, various relationships between the tiles change, things that are knowable at time t become unknowable at time t+1 etc. The way the cellular automation is described now is just a toy example to illustrate the basic structure of discovery space: that is, how what is knowable is channeled and contoured by the unknowable and how the known reveals the unknown.

This simplified model of exploratory behavior, in which foregrounded knowledge is made clear by contrast with its background of ignorance can be likened to the ancient Greek concept of *alethia* resurrected by Heidegger. It is related to Heidegger's concept of *Erschlossenheit,* world-disclosure. Alethia is sometimes translated as truth, but this is a gross underrepresentation of its nuanced meaning. In Heidegger's parlance, it means unclosedness or unconcealedness. Not the truth, but what reveals the possibility for truth. Rather, it designates a situation where the truth is knowable. A decent approximation for it would be the word "discoverability." If A implies B, knowledge of A discloses B. Knowing the laws of electricity, discloses the possibility for knowing electromagnetism. More generally, the nature of the world as it is perceptible to us discloses the possibility for knowing it. To become knowable is to first become intelligible, to be open for access.

In the essay *The End of Philosophy and the Task of Thinking* Heidegger uses the metaphor of a dense forest in which no light can reach through the canopy to the forest floor. Eventually, however, one comes across a clearing (*Lichtung*) an open space where there are no trees. Inside this forest clearing, the truth becomes intelligible, the light of reason manages to shine through to the bottom. "Light can stream into the clearing, into its openness and let brightness play with the darkness in it. But light never first creates openness. Rather, light presupposes openness." Such a clearing discloses the world, it is an opening that is "not surrounded by what is; rather, the lighting center itself circles all that is, like the Nothing which we scarcely know." (On the Origin of the Work of Art).

Note that Heidegger states that both darkness and light play in the clearing. That is to say, both knowledge and ignorance as absence of knowledge, are present in the clearing, the clearing has both positive content and "present absences" that are apprehensible. The dense forest in contrast has no presence at all, it is inscrutable to reason; thus even the possibility of knowing it or not knowing it is undefined. It is where reason cannot go. The analogy is quite similar to the concept of cognitive closure—we can only think that which the fixed architecture of the mind makes thinkable. Heidegger's concepts of world disclosure and clearing aligns revealingly with the concept of the structuring

relationships between knowledge and ignorance. In an exactly similar way knowledge, discoverability and ignorance stand in relation to one another. Alethia discloses the possibility of both what we can know and are ignorant of. A *deeper*, absolute ignorance meanwhile surrounds and encloses this clearing. This is the sort of ignorance that cannot even be conceptualized, that falls outside the scope of reason itself.

What is essential for human purposes is the openness, the discoverable truths that the clearing within this absolute, impenetrable density makes manifest. "The clearing is the open region for everything that becomes present and absent," a strict limit. And "…the light of reason throws light only on openess, it does not concern the opening, but so little does it form it that it needs it in order to be able to illuminate what is present in the opening." The opening, in other words, is the modal disposition of knowledge, it is the the knowability of certain facts that are disclosed.

In a sense our own clearing in, absolute terms, is our physical light cones, although we certainly have many practical and compartmentalized clearings within this sub-dividable upper limit. We can't know how much and what exactly is going on beyond it. But we have our own subjective light cones too, our own horizons of discoverability. In Heidegger's characteristically enigmatic phrasing, "…what is absent, too, cannot be as such unless it presences in the free space of the opening…The path of thinking..needs the traversable opening…in that opening rests possible radiance, that is the possible presencing of presence itself." The discovery of an absence, of an occluded ignorance, is a byproduct of its having been revealed in the opening, a knowable unknown. We *notice* its occlusion. It is discovered by its contrast with the light streaming around it, the field of reason that envelopes it. Anything potentially thinkable must be present in the clearing. The absence is *present* as an absence through its disclosure within the unconcealed clearing of cognitive possibilities.

At this point it should be clear how this Heideggerian digression feeds back into the general line of thinking. World disclosure is another term for the conditions of possibility of knowledge, a certain "grey", remediable type of ignorance, an opening in a free space where knowledge can be obtained. This is

not the hard, insoluble "dark" ignorance of the dense forest, that region outside our subjective light cone or outside our objective light cone that cannot be accessed, which we not only don't know, but which is never presented as a potential candidate for knowledge in the first place. The clearing aligns with the boundary of cognitive closure; beyond it lies an impenetrable nothingness— indeed it is *not even* a nothingness— it cannot be named because it lacks accessible presence by which we can assign a label to an object. This surrounding oblivion nevertheless functions as a shape-giving edge. Just as absences, as the darkness playing against the light, are given presence within the disclosure of the clearing, so does the dense forest of negative space beyond the clearing lend shape to the clearing itself.

World disclosure is based on the concept that phenomenologically the world is never presented as one salient, completely accessible and intelligible totality. The unconcealment (alethia) of the world is not the same as knowing the world, it's knowing that you can know something about it. It is a privative. But what exactly is it depriving of? Concealedness is for Heidegger not a mere state of ignorance, but the negation of any epistemic value whatsoever. To unhide is to remove what removes, not to know but to make knowable.

Another facet of Heidegger's philosophy is useful here. These are the concepts of throwness (Geworfenheit) and projection (Entwurf). In the sense Heidegger originally intended, thrownness represented the state of being "launched" into existence at a particular place and time, born into a world at a given moment without context or explanation, and then being forced to get one's bearings. But there is another, epistemic sense that this term can be used. In the face of the unknown we are also thrown, having to piece what's happening together minute by minute, hour by hour, day by day… Eventually after much trial and error and struggle, if we are lucky and smart, we reach a plateau of stability in which for the most part our world in broad outlines has become navigable. When we reach for something, we find it there, because we have fleshed out our concept of the world by building our model of it. The unexpected bolts down like lightening from time to time, sometimes too often, but for the most part we have acquired a crystalized invariant representation of our

epistemic context that allows us to conduct ourselves with a modicum of frictionless fluidity.

Projection, in Heidegger's parlance, speaks to the state of being in which Dasein (the experience of being human) reconciles itself with the world and asserts its possibilities over Being. In other words, it concerns the agentive aspects of human experience that visualize goals and assert those actions which actualize potentials. That is to say, we are not "cognitive couch potatoes" in Andy Clark's turn of phrase, passively awaiting inputs from the external world to make cognitive determinations. We are seekers; we are goal-directed intentional agents. Our existence is made up of countless tiny guesses which we then act out to confirm or disconfirm and proceed to build on. Projection then is also the will to make what we wish to be real true, to forge our ideas into realities. The resolution to the predicament posed by thrownness is then to project our aims onto the world, to meet it in middle and to then diminish the difference between our projection and the reality until they are identical.

"Projection" is used in another related and equally relevant sense to mean prediction, forecast, or expectation. The predictive processing model of cognition states that the main task of the brain is to construct a representation of the world which anticipates and exploits its salient features, constantly comparing anticipated output signals with oncoming sensory input signals. Inasmuch as we are unaware of the difference between our expectations and the true outcome, or unable to establish grounds for expectation in the first place, we are ignorant. To the extent that our expectations are contradicted by the stimulus, we learn. Generally such attempts to match expectation to reality take the form of induction.

The problem of induction is notorious and longstanding. How can we generalize from an instance to a class of objects, having never— and usually never being able to—observe every member of the class? The problem seems to signify an essential contradiction at the heart of empiricism. Observation yields empirical evidence, and yet we never acquire comprehensive observational coverage for any class of phenomena. We trust that the partial observations we have made speak for the unobserved. To test this, we marginally support our

observations by making a new observation and confirming the assumption. But this is only ever an incremental source of support and it never brings us to complete certainty. At any point the next observation could buck the trend. Hence, black swans may lurk in our midst unbeknownst to us, and nothing in our observational database saying all swans are white can provision certainty that there are no black swans. Induction, therefore, provides us only with a method for judging the probable, but not for ruling out the impossible. And so long as the possible still remains, the improbable no matter how rare still lingers in the shadows and could detonate the best laid plans with barely a moment's notice.

The ghost of ignorance haunts induction and even our most storied knowledge producing mechanisms of science cannot escape it. So the question stands before us: what to make of induction in the scheme of things? How does it unfold in relation to ignorance?

A rule of thumb is that the simple a fact is the easier and more likely it is to know. The more complex it is, the more likely there is to be ignorance about all or some portion of it. The reason for this is plain: a simple fact contains no relations to other facts that are prerequisite to its ascertainment, while a complex fact consists of several facts and their relations. The more complex an object the more degrees of freedom it has, the more possibilities it is capable of exhibiting. People are hard to predict because people are complex. By this rule then, a complex object that displays predictable results must not be a complex object. Complexity begets complexity.This straightforward observation reveals a valuable truth about the shape of discovery space.

A simple fact is easily definable. It is a fact which does not require other facts to be understood. A complex fact is a structure of related facts, it is a system of facts which must bee comprehended both in terms of its individual parts and as a whole to be fully known.

The version of epistemology discovered here is comfortable with a probabilistic view of things, particularly as it relates the known to the unknown. As a necessarily incomplete relation, the known, if it can contextualize the

known unknown at all beyond the brute fact that it is unknown, grasps it only as shades of possibilities that might be differently weighed in terms of their relative likelihood. Background knowledge may not be sufficient to directly infer the unknown, but in certain cases it can justify speculating with a fair degree of accuracy about the probable subspace of the unknown's possible dispositions. A simple unknown reflects a simple fact, a datum that can be immediately apprehended and does not require the additional discovery of various interdependent facts and which by itself does not afford overly numerous possibilities. This leads us to an amended vision of discovery space.

Discovery space is far from the two dimensional, flat sort of layout depicted in the previous section. It can also be given a "third dimension" in the form of a hill climbing dynamic. The more difficult a discovery is, the more complexity it holds, the steeper the hill and the higher the peak. Accordingly, the more effort that must be invested climbing the hill to get to what's behind it, and discover facts that have the first fact as a prerequisite. Simple facts in contrast are barely bumps in the road, whereas known facts and trivial facts are perfectly level. 1+1 is easily learned; differential equations not so much. Using the measure of Kolmogorov, or descriptive, complexity, we can infer how the probability of ignorance is a function of the shortest possible description needed to characterize a phenomenon rather than the simple binary condition of knowing or not knowing it. What is more probable to know are those facts with concise descriptive complexity, while more complex facts, like how the economy works, has no obvious shorthand synopsis.

Descriptive complexity defines the absolute minimum amount of information needed to translate an object into an algorithm that could produce the object as output. Alternatively it defines how much description it would take to fully characterize a system's components and behaviors. By this measure those entities with lower descriptive cost are easier to know and predictably more commonly known. Knowledge by description is one well known type of knowing, and as such it fits into this formula neatly. A satisfactory proof of knowledge is the capacity to describe the object to completion, in such detail that the description comprises a "recipe" for the object. As such, the more complex the object, the

greater the demands on description, the higher the learning cost invested, the less probable it is to be known.

This general notion of complexity suggests that some facts are inherently harder to know while others are more like freebies. Thus the distribution landscape of ignorance more or less maps to a landscape of complexity . Information complexity inhibits the immediate reception of certain complicated systems of facts. As a heuristic, we would then anticipate that the more complex a subject the more likely people are to be ignorant of it from a population statistics standpoint, which not for nothing explains how specialization and expertise is skewed in the economy. If everyone could easily apprehend the requirements to be a brain surgeon, there would be no need for such surgeons to subject themselves to arduous years of medical schooling and the painstaking acquirement of precise technical skills, or for them to command impressive six figure salaries. Theoretical physicists would be born rather than made. Newton would be mediocre. What we call a career is the lifelong dedication to climbing a particular incline in a mountainous region discovery space.

Complexity and ignorance are deeply interrelated. While it is certainly possible to have one without the other, their connection is a priori bound to the fact that the more working parts in play for a system the lower the joint probability of knowing all of them compared to the probability of knowing any single one of the facts in isolation. Uncertainties, when compiled, do not cancel each other out but amplify the overall uncertainty. The much is simply intuitive.

From another angle it can be said that anything which becomes known must be simple enough to be understood even if it is not simple in an elementary sense. Science, though Occam's razor, and mathematics, through the procedure of simplifying multi-termed expressions into shorter expressions, both stress the value of simplification. Similarly, our knowledge representations are not stored as ponderous lists of details but hierarchical schemas that organize related concepts under shared headers. In the course of solving the problem it is invariably simplified until it is in a form that falls within the channel capacity of the mind. Conversely the more complex a phenomenon, or the more complex the level of description at which it is analyzed, the more it approaches the

unthinkable. Consider statistical mechanics. The entire premise of the field is that the local behaviors of the molecules are stochastic and unpredictable. However when viewed as a large ensemble consisting of very many particles the system exhibits lawful regularities. Thus an order of complexity is shed so that the system is defined at the scale which describes its aggregate behavior as conformable to a rule . In this case a sweet spot is reached where the simplest form of description is also the one that captures the most information about how the system as a whole behaves, but this is only because e got lucky. Attempting to develop a physics of gases on the level of description that models each particle's individual behavior is computationally intractable for all but perhaps the beefiest of supercomputers and likely beyond the scope of individual human understanding because it requires the modeling of many millions of random variables independently. Regardless of comprehensibility, such an overwrought model of statistical physics is a terrible format to encode knowledge within.

The brain may be powerful but it is also bound by striking constraints. In the classic paper The *Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information*, psychologist George Miller found that human working memory capacity was on average limited to seven or so distinct contents at any given time. ([http://psychclassics.yorku.ca/Miller/]). By applying principles of information theory to measure memory processing in terms of bits, Miller's data is in line with the relationship between complexity and ignorance described here. "When we have a large variance, we are very ignorant about what is going to happen. If we are very ignorant, then when we make the observation it gives us a lot of information. On the other hand, if the variance is very small, we know in advance how our observation must come out, so we get little information from making the observation," Miller states. Complex systems thus display a tremendous amount of variance and so their limited comprehensibility is a function of how much information they contain and thus how much burden they place on memory. Miller considers what he calls multidimensional stimuli, aggregates of independently variable attributes of a particular sensory configuration. As variables are added to the display, the accuracy of the test subject's responses decreases, approaching an asymptote

Miller denoted as the “channel capacity.” Once the added stimulus components exceed this channel capacity the accuracy of judgments falls precipitously. His results led him to the number seven as the mean amount of subitziable information units that could be held in working memory at once.

So the low-hanging fruit, those fact which contain minimal variance, tend to be better known while those showing a greater degree of variance, complexity, less so.

The flip side of the coin is what has been called “cognitive miserliness” in the psychology literature (Stanovich, Keith E. (2009). "The cognitive miser: ways to avoid thinking". What intelligence tests miss: the psychology of rational thought. New Haven: Yale University Press. ISBN 9780300123852. OCLC 216936066). Generally people seek to avoid unnecessary cognitive effort and tend to prefer interpretations which follow the “path of least resistance” and fit their prior understanding rather than to struggle to bore through the dense mantle of stubborn conceptual difficulty. Processing fluency, the ease with which a problem is understood, goes a long way to determining what exactly we end up understanding. Thus our own cognitive bottlenecks have a tendency to produce ignorance as a byproduct by favoring what is easier to understand at the exclusion of what is harder, leading us to risk construing what is easy with what is true and what is hard the false.

What one might be tempted to call knowledge is often the locally accessible undershot of a broader and more comprehensive yet impenetrable truth, the appreciation of which demands an intolerably immense investment of effort. So to know, sometimes, is to settle for what can be known within constraints rather than what is to be known. Consequently what we can understand is mistaken for the true state of things in all its richness; it is instead a diluted and fragmentary abridgment of it. This does not imply that what is gleaned from the towering object of the unknown is false because it is incomplete. Rather it is imperfect knowledge not because it is false on its own terms but because it is falsely believed to represent the whole picture.

It is partly true that “complexity is in the eye of the beholder” and that what is simple to a higher or more studied intelligence is more complicated for a lower

or less educated intelligence. That's not the whole story however. It does not change the basic reality that certain problems are denser and more convoluted than others by nature in terms of their descriptive requirements. In means only that a given agent might have a better affinity for simplifying that complexity than one which is unable to gain purchase into its apparently inscrutable and indistinct edifice thanks to education or some other differential. Those who struggle to understand it are unable to perceive cracks along its surface that might be taken for "cut here" marks that would separate the phenomenon into simpler, modular distinctions that divide the problem into more manageable chunks. They instead only see a smooth, uniform and impenetrable bulk about which nothing can be done and no distinctions can be made. As is well understood in computing, the greater the resources in processing power and memory, the faster a problem is solved. The complexity of the problem is totally independent of the quality of the equipment deployed to solve it. Intelligence by one definition is then the capacity to simplify the unknown to a level that makes it understandable; it is greater complexity overpowering and subsuming lesser complexity by reducing it to simplicity. Lack of intelligence in contrast is a simpleness that, by virtue of its lack of complexity, fails to transform a complexity greater than it to a format simple enough for it to understand.

One can conjecture that there is a cognitive bias against complexity or in favor of simplicity. There is no a priori reason that most parsimonious conclusion is the right one, or that the fundamental equations of physics must be elegant, or that truth is beauty and beauty truth. Occam's razor, which holds that the theory which takes on the fewest assumptions is preferable to one that assumes more, is itself not necessitated by any empirical observation. As evolutionary psychologist John Tooby pointed out, human beings seem to have a natural tendency to prioritize a single essential cause when explaining something even though any complex phenomenon should be expected to have multifactorial causes([https://www.edge.org/response-detail/27168]). A sitting president is often blamed for everything that goes wrong with a country during their tenure, as if they single handedly run the economy and have omnipotent control that they simply failed to execute on. Such reductive thinking is a byproduct of the

challenge complexity poses to knowledge and the need for an answer in circumstances that are too perplexing to sum up in a condensed scheme.

In *Discourse on Metaphysics* Leibniz claims, "it is not not even possible to conceive of events which are not regular." Quite a strong statement, but it hints at a basic truth. To conceive of an event is to conceive of it in a format that the mind accepts as a frame. "Our minds," Tooby writes, "evolved to represent situations in a way that highlighted the element in the nexus \[of causality] that we could manipulate to bring about a favored outcome. Elements in the situation that remained stable and that we could not change (like gravity or human nature) were left out of the representations of cause ([https://www.edge.org/response-detail/11211]). Such emphasis on key causes rather than nexus causality has its shortcomings. Such truths as we do apprehend necessarily entail a degree of misrepresentation, an insensitivity to the over-arching nexus of causality.

Contrary to notion that biological fitness should maximize representational accuracy, so that the senses give you the truth about the external world and the mind is a truth machine, biology instead only surfaces what's relevant to the organism's lifecycle and manipulates how it is perceived to favor the organism's survivability. One would think, contra this hypothesis, that a cognitive system that is presented the highest fidelity model of the external world would outperform one that receives a distorted picture of it, but this misunderstands the evolutionary role of perception. Does the prey have any objective reason to prefer its own life over the predator's meal? Bees see ultraviolet light, which is invisible to the human eye, because it benefits their fitness, not because ultraviolet light is a defining attribute of flowers. If Hoffman (2018) (or by degrees Kant) is to be believed, perception is like a virtual reality display that superimposes an internal schema on external affairs and selectively draws out salient aspects of the environment, adding "coloration" to the otherwise neutral information outlay to highlight certain goal-relevant facets. Biology can therefore attribute properties to the sensory object that do not inhere to it. Even color vision is an embellishment in this respect. Colors are purely a function of the visual system and are not bound properties inextricably attached to the perceived object.

The key point in all of this is that we are always, deceptively, on the lookout

for key points. There is no rational necessity that urges this impulse to reduce complex phenomena to their central tendencies other than that it renders the problem thinkable. Even our valid understandings, therefore, might be tinged by a subtle blemish of ignorance, as they may leave out an arbitrary amount of extraneous detail which is equally descriptive of the problem in all its fullness. Evidently people are still able to solve problems while subject to these distortions, but that is because we understand the issue in a way that is narrowly relevant to us rather than in a purely objective sense. The same inclination however does harm as well as good.

Framing conceptions in this way usually spells a residual byproduct of ignorance, because the frame excludes as much as it includes. The most accessible features of the filtered information are in turn passed on and filtered until they are reduced to their most elementary explananda. Incoming signals are squeezed into a shape that facilitates processing fluency. The various "determinisms" that have marked the history of ideas, from the economic determinism of Marx, the psychic determinism of Freud, or the hereditary determinism of Spencer, are remembered today mostly for their limitations rather than for their profundities. Regardless whether the express content of these theories have any validity, they are not and cannot be the "whole story." They are the products of a simplification process that shoves complexity into a comprehensible mold and discards whatever overflow fails to fit it. Generally it is a good idea to capture the most fundamental aspects of a causal process rather than describe it as an uncompressed, hard to follow, diffuse volume of listed details, but let it be stressed that this reflects the limitations of the mind and not the reality of the observed phenomenon.

The final consequence is that we should expect the distribution of knowledge to negatively correlate with the distribution of complexity on average. The more working parts a system has, the more there is to know about it, the less intuitive what there is to learn, the more likely that the thing belongs to ignorance. What's worse, the general tendency to reduce phenomena to their central causes is not dictated by any natural necessity but rather arises from the cognitive constraints of the human mind. To think about any part of the world one must exclude from

consideration the rest of it, or even parts of the object that seem less explanatorily relevant. Attempts to explain too much with too little, so called greedy reductionism, however tempting, ought to therefore be resisted. Even the cleverest arguments may obfuscate complexity and hide a store of ignorance.

The Problem of Induction from an Ignorance First Perspective

The problem of induction is a longstanding conundrum in epistemology. It is, if not the central problem, one of the hardest to understand, and the most important. Without a means to strictly specify how we acquire valid knowledge about the external world and its patterns, we are forever in the dark, fenced in by limited information, reduced to petty guesses and doomed to unendingly violated expectations. The primary attribute of induction is that it is imperfect, probabilistic. It cannot be raised to the perfect standard of deductive inference, in which perfect information conditions allow the entire situation to be evaluated and no relevant fact is hidden.

It is probably not possible to develop a better method of induction, although scientific induction represents a refinement over and above intuitive inference. That's because the conditions in which induction unfolds are imposed by *structural ignorance*. Induction is primarily hedged into its flawed position because ignorance crowds out any possibility of a deductive information landscape.

Induction is in large part a function of temporal ignorance—it is mainly future-facing and regards how past and current information might predict the future. It is also about category formation, the problem of how to classify objects that bear resemblances, grouped under a prototypal schema that captures their shared definitive properties. The argument can be made that these two different aspects are one and the same. They are united by their relation to ignorance which expresses itself both in the formation of categories and in time but is invariant regardless of circumstance. Both the categorical and the temporal subproblems of induction are consequences of a single cause: structural

information incompleteness.

The problem of induction is also characterized by the fact that the universe is interesting. In a sense, the problem of induction is attributable not to the finitude and incompleteness of our knowledge, but the dynamical nonlinear nature of the universe itself. Its causal structure affords many possible instantiations and configuration states for its objects. If everything happened the same way for eternity all inductions would be perfect. Induction is high stakes because the outcomes it quantifies over are nontrivial.

One must stop to wonder that if the *universe is undecided and incomplete* that inferences about it must necessarily be as well. It goes back to William James' suggestion of the unfinished universe. No wonder the inductive inferences continually fall short and are restricted to mere probability. The universe itself is pending, filling out, unfolding into its leading edge and expanding its horizons. By inducing the future we are probing its frontiers. The present is a "spotlight" a briefly actualized snapshot of the universe's information state. It constitutes a single "build step" in the course of the execution of the universal algorithm as it expresses its gradual actualization in time.

Inferences can only be extended to facts. So it stands to reason that if future events do not exist except as potentials, the only effective inferences that can be made are not about the future events themselves, *but the potential for them to happen in the present*. But the key fact about a potential is that it only might actualize, it is not necessary that it must. A planted seed has the potential to become a tree, given the right conditions. A drought would put an end to that. So it is for potential future outcomes in general. The temporal mystery of induction is eliminated once we understand that it reasons over potentials in the present, and those potentials are the result of past causal history structuring the world into a given state where energies and objects are at the precipice of a near-term future outcome.

Uncertainty can be conceived of as the recognition of an array of conceivable conflicting possibilities, the outcome of each of which is dependent on the

satisfaction of given conditional requirements which are wanting. Uncertainty is the recognition that no obvious outcome is apparent. Probability is quantified possibility. The probability of possibilities is considered to be function either of how frequently they have happened in the fast or how consistent they are with prior beliefs. In some way the entire event structure of the world can be be represented as a vast decision tree of conditional statements, as "if-then" yes or no questions that logically mirror causality and the flow of events. The problem of induction concerns the conditions which make it impossible to locate the necessity of any one if-then causal linkage, to assert "if and only if then."

Of all forms of inference, induction is the one which faces ignorance most directly. It is the if-then that must be answered at the edge of experience at the edge of time, poised at the point where the past and present information is abruptly cut off by the blankness of a severed and informationally noncommunicating future. One would be wrong to use deduction in such information incomplete settings. Deduction is the logically stronger form of inference because it operates inexorably and quantifies over an information complete environment. If the right rule is applied to the right case, deductive reasoning is never false.

One can be wrong with induction because complete information required to answer correctly with certainty is strictly inaccessible owing to the immutable event structure and dynamic causality of the physical universe. The information to know that *this* necessary outcome is not only inaccessible, but perhaps nonexistent. This may be because, for all intents and purposes the future one is inferring does not exist yet, so facts about it cannot be asserted. Induction seeks to reach past the boundary of an information limit and to characterize the absent by the present, the unobserved by the observed, the actual by the potential. In the battle with ignorance induction is on the frontlines.

Any important decision usually comes at some risk. What is risk? Risk is the possibility of being wrong and that you will be penalized for it. Risk is a fact of life bounded by ignorance and incomplete knowledge. Had we magical and perfect insight, there would be no risk, only actions you should either take or avoid. One would simply select a goal and make the right decision infallibly.

What company would hire me? No need to play the field, simply apply and speak to them. What would make me rich? Follow any given recipe for action. In such a cognitively godlike state there would be no choices: we'd see the one or all of the pathways through the network of possibility that leads to our goal and take it. It would be like following a string through to the center of a maze.

Yet we live in a world that imposes uncertainty on us and as such structural ignorance is ubiquitous. Usually whatever is certain is what carries minimal importance and therefore risk, it is low stakes, and it requires little thought. You don't have a problem if you know the solution. A problem is always prompted by ignorance, by uncertainty. Generally if we are apprehensive about something we are ambiguous as to the outcome. There may be cases where we are anxiously awaiting a certain outcome, like a guilty prisoner awaiting their inevitable sentencing or an ill patient awaiting their diagnosis. For the most part people have a terrible reaction to uncertainty. As a matter of fact uncertainty is considered a "major predictor of stress" in hospitalized patients awaiting treatment as one might expect ([https://onlinelibrary.wiley.com/doi/abs/10.1002/nur.4770070304]). Intolerance of uncertainty is a surefire human trait.

Wherever there is uncertainty there is a threat to the integrity and wholeness of our knowledge, and therefore the stability of our mental world. The important thing however is not to run from uncertainty, but to embrace or even utilize it.

One way to characterize the universe is as a random sequence of high entropy states, with repetitious, ordered local substates that variously blend into or interact with the surrounding disorder. Such substates and superstates flow into and out of each other through the exchange of energy. Earth, compared to the surrounding vacuum where nebulae and cosmic rays boll around haphazardly, is an example of a localized complex state contrasted by global chaos. In a manner of speaking, earth represents knowledge, while the void of space ignorance from the perspective of information space.

Statistical mechanics describes change as a diffusion of thermal energy from higher to lower concentration gradients expressed through heat dissipation approaching the equilibrium. The universe at large is dissipating energy until it reaches a "cosmic room temperature." Along the way at the longest timescales,

information entropy increases because the universe gets more random. The only thing certain about randomness is that it *is random*, it is perfectly predictable in its unpredictability, but randomness is *usually* uninteresting. An infinite randomness contains all organizations that can be formed of its base elements, but none of them can be usefully accessed. An infinitely long strand of random DNA would contain the DNA of all organisms as a substring. If the sequence wasn't random, an inductive rule could be derived to demonstrate that there are some genomes that aren't included, because there would be a limited number of possible subsequences in accordance with the rule. Randomness in itself cannot be efficiently "mined" for such useful patterns tucked away in it because the energy cost to detect patterns within it is too variable. One might have to search forever to find a pattern that matches one is looking up. The desired pattern may appear once in one spot, and then not again for another million indexes, and then again for another quadrillion five hundred and three, or never again.

Randomness is certain in its uncertainty. Randomness, like ignorance which is its correlate, still has structure despite this, as one can infer properties of randomness. Though a sufficiently expansive randomness may contain useful possibilities, we lack tools by which to isolate them. Randomness is not related to possibility without some generative mechanism that "refines" the randomness into a more ordered state which accesses its hidden structures. Evolution is one example. An essentially random sequence can be refined until it becomes more particular and serialized in its arrangement. Randomness is an anti-pattern defined by the fact that it is not characterized by any of sub-patterns. Randomness cannot be compressed. And to compress, to simplify, is to know. Thus randomness is the mathematical correlate of ignorance.

To describe randomness , to determine that a sequence is random, you have to describe the entire sequence and at no point find a repetitive subsequence which can be used as a means of induction to characterize the remaining dataset. A perfectly ordered sequence, in contrast, is describable by breaking it down into its distinctly repetitious subsequences. Both cases of perfect randomness and order are characterized by uninteresting triviality. They are the limiting cases, the polar extremes, about which we can be certain of their absolute noisy

randomness or static order . Both are uninformative. It is somewhere in the middle where things get interesting. There, you have systems such as ourselves and life, which have just the right amount of order to maintain stability through time but also the right amount of spontaneity to permit evolutionary change.

Laid out this way, the accuracy of induction is a response to the measure of complexity inherent to a system, that is, it is measure how many possibilities are afforded by it and how . Induction is valid in situations that are dynamical enough to be occasionally surprising but redundant enough for the past to speak for the future at a frequency greater than chance. This is the state in which everything important happens. An entirely random phenomenon cannot be induced.

The analysis of induction tends to be focused on its logical and cognitive properties rather than the external physical and informational conditions that constrain and channel those properties and impose their limitations. Induction is something we can't help but do, logical flaws and all, because it is one of the few means by which we can attempt to cut through the fog of an uncertain world. In a manner of speaking, the fog is always there even though we cut through it, it is never dispelled, once it is penetrated fog moves in to fill the gap again. The problem arises when epistemology is only viewed solely in terms of its positivistic half rather than holistically, including a proper understanding of the structuring effects of ignorance on epistemic states. Induction does not make sense when framed in terms of how finite knowledge is extended from itself given an incomplete sample of observations. It is better understood from an outside-in perspective that appreciates what's going on outside that knowledge to condition it.

Why exactly is induction considered problematic? According to Hume's famous treatment of it the problem of induction consists of two parts.

1. Generalizing from some observed members of a class to all members
2. Inferring that the future will be like the past

On close inspection, the second part is the application of the first problem to the dimension of time. It is the problem of classifying past, present, and future moments as members of the same category of causally connected events. So

really the induction problem is the issue of how to make valid generalizations from partial experience to all experience or at least over some quantified universe of discourse. Induction differs in the same respects from deduction but in reverse: with deduction we generalize from the class to the member, because we have the class already fully formed and specified axiomatically. And with deduction there is no a priori concept of time. Which is to say, there is no causality within the symbol space other than its strict rule logic; everything is given immediately as an information complete environment and it is only our thoughts about it that must unfold over time as we pass over the symbols. The introduction of time and with it incomplete information, is one thing that distinguishes induction from deduction.

With induction we must go from the member to the category because we experience time as broken up into an infinitesimal series of time slices, and so we must knit them together piece by piece to reconstitute the idea of causality. To compound the problem each time slice contains only a sliver of information about the whole. For the same reason we never see the future. We are structurally ignorant of it because the laws of (classical) physics determine that information can only be transmitted in certain ways and the future cannot communicate with the past. Physics entails what is called *time asymmetry*. Thermodynamic entropy dictates that the past was in a more ordered state than the future and therefore processes in time are irreversible because the future is incongruous with the state of the past. To relate the past to the future and to affirm the uniformity of nature (the universal applicability of natural law) is to generalize from the past, or rather, to identify regularities in the past that are expected to remain constant irrespective of any other changes in the spatiotemporal distribution of objects. Such an invariant framework, which specifies the possible recombinations of objects, is called the laws of nature. They are the special class of properties which “cut nature at its joints” and define its most fundamental rulesets (Loewer, 2007).

Chapter 7 claimed that one of the few, but also most insightful, certainties or more accurately minimized uncertainties is natural law. But it would seem that if the problem of induction is a concern then the integrity of natural law slips into

doubt. Natural laws are the ultimate product of induction. They do not however constitute the logical rules of the axiomatic system of nature which can be used to derive any physical instance. Rather only can the observation of the instance be related to the natural law. From an understanding of the natural laws alone without reference to particular physical facts, we could not apply a solution to every problem in the physical world. We just know that the solution, like the problem, must be consistent with it.

Natural laws are para-empirical statements. There is nothing in the world one can ostensively point to which could be taken as a token representative of natural law. Just as Hume pointed out causality in itself is not empirically observable but only inferred from from recurring spatiotemporal relationships between empirical objects. Natural laws are not "out there", an empirical object among objects. They are a summary of observational regularities and as such they must drop all specific deals and retain only the invariances across heterogenous instances. An apple falling from a tree onto Newton's head is not a natural law, yet it is still an example of it. Natural laws are not empirical. They are generalizations over sets of observations but do not belong to the propositional content of any single observation statement. Rather, they are the rule that characterizes a structured (if-then) change in all the observations given certain stipulated conditions present in each of them.

(Inquiry into human understanding Hume)

Natural laws are important for the theory of ignorance because they place bounds on empirically *meaningful* ignorance. If we are willing to accept that natural laws are true and valid ideas, then they extensively economize empirical space, they scaffold ignorance in a rational frame that subordinates all material facts. Just as there may be limits to knowledge, there must be limits to empirical ignorance. Natural law specifies exactly those restrictions. As mentioned in 7.1, such universal generalizations permit reason to rule out certain possibilities and confine predictions to a bandwidth of plausibility. A distinction can be made between *junk* and *valuable* ignorance-as-potential-knowledge, that is, truths worth knowing versus ones that aren't. Various unknowns, just like various knowns, have different valuations in terms of practical importance. The internet

phrase, TMI, "too much information" casually hints at this distinction, there are some things one is better off not learning of. True wisdom and useful potential knowledge is sprinkled throughout a wasteland of trivialities. More than this, there are numerous junk statements, gibberish, falsehoods, and farfetched counterfactuals that one is technically ignorant of in the sense that one is not experientially acquainted with that are simply not worth knowing.

Do you know what the billionth digit of pi is? That may be worth knowing in the very limited case of a very specific calculation, but even this example alone shows that when it comes to trivia such as knowing the nth digit of pi, there is an actual mathematical infinity of things one doesn't know because pi is an irrational number. What is not worth knowing may vary depending on one's purposes, and indeed there may be an irreducibly subjective element to it, but in general nobody needs to know that one space rock fragment lightyears away broke off from a bigger space rock, or that a particular ameba split in two just now in a particular pond in a swamp in Louisiana.

Yet precisely what is not worth knowing subtractively defines what is worth knowing. The laws of physics, as the crown jewels of induction, allow us to divide the superstitious from the plausible. Other valuable, accurate generalizations do the same, but these must be consistent somehow with the laws of physics.

Induction is critical to the understanding of ignorance because it is pivotal to the cognitive process of category formation (Rakison, 2007 [https://www.cmu.edu/dietrich/psychology/infant-cognition-lab/publications/Rakison%20(2007)_2.pdf]). We formulate and refine our perceptual categories through a Bayesian process of inductive reinforcement beginning in early childhood. Each occurrent observation readjusts or confirms the category to which its object belongs. If no previous category exists, given time and enough repeated overlapping experience, it will be induced. How it is induced, is difficult for us to understand because it happens unconsciously. Nevertheless this category induction is essential to the understanding of how the mind copes with its ignorance of uncountably many particulars by constructing generic classes that abstract the shared traits of each member of a category.

We can never experience each tree in the world, closely inspect the fine details of its brachiation and the patterning of its bark or count every leaf. We may not even have the words to describe how the details of two trees differ, even though on close analysis every tree, like a fingerprint, is unique. No two trees branch out in precisely the same way, or have precisely the same count of leaves. Yet none of this matters to us: so for each tree there is the category Tree.

What we can have is a generic knowledge structure of trees, and if we are particularly passionate dendrologists, yet more distinct subcategories of each genus and subtype. Ultimately, upon sighting an individual tree, we can then unpack and apply this categorical knowledge to it individually. To store our knowledge at the lowest level of particulars is seriously inefficient. Doing so would never tame the world's multiplicities. What we do is categorize, amalgamate, and abstract. And categorization is the product of induction.

This economization of thought by induction is nothing short of an adaptive strategy for coping with information overabundance and the ignorance that it imposes on our memory and related cognitive constraints. To generalize is to be ignorant and knowledgeable at the same time. The type is privileged over the token, and for good reason. Without doing so it would be impossible to compactly represent the boundless pluralities of objects constituting the world.

This sort of inductive categorization is not flawless. In some ways it papers over the cracks of the gaps in knowledge. It *invites its own sources of ignorance* directly associated with the problem of induction, in the course of alleviating other ignorance. The fallacy of hasty generalization, in which a broad conclusion is made on the basis of too little evidence, is just one such example. But this is merely a failed inductive generalization. Like a basket ball throw that misses the net, it does not mean you can't get a basketball in the net. Ironically, the very arguments *against* induction are inductive! How do we know when to halt? What in the information stream informs us that we have collected sufficient information to sanction a valid generalization? The simple answer is that we don't, that nothing informs us of such a step within the information stream itself, and this is where Hume rests his case. This is the dilemma of the inductive leap which been a perpetual source of confusion. By what grounds can we assert the

validity of an inductive inference beyond that of faith or merely subjective confidence? As Popper pointed out, a principle of induction is needed to certify induction's validity, but this self-same principle cannot be induced through experience (Popper, 1959). No experience provides the basis for asserting the validity of induction.

It is also possible to over-generalize, to often fatal effect. As when someone stereotypes a person's behavior based on the characteristics they share with a group. Or when we cannot see the "trees for the forest" and fail to discriminate the idiosyncratic features of a situation that is similar but different in crucial respects from the class to which it belongs.

According to Popper scientists, and presumably by extension reasoners in general, cannot go from instances of observation to a theory quantifying over every possible observation of that class. This would assume some missing link in the singular statements that connect them all to the universal statement. Nothing inside the mechanism of observation, in other words, provisions the inferential equipment necessary to carry out theorization about the class of phenomena. In a similar way Hume supposes that we do not infer causality from objects, rather, we intuit connections between our experiences which form a separate system of impressions from the objects themselves (Inquiry into Human Understanding).

The mystery of the inductive leap disappears when we consider that the inductive reasoner is not at fault for their incompleteness: they are fated to uncertainty because the very order of the world is fuzzy and unfinished. Induction cannot be perfect, not because of a flaw in human understanding,* but because of a deficit in the world's legibility*. If the world were a text, it would be like an old book that is weathered and moldered. Some parts can be made out but others are unreadable. Sometimes the unreadable parts can be inferred from the readable parts, but there are also stretches where no such continuity can be discerned. The pages of the past are faded, missing, or frayed. And not only that but the entire sequel to the book, The Future, hasn't even been written, even though we're fully aware it comprises an all-important part of the story and will fill in all the plot holes left open by the first book. Is the reader at fault for having a hard time understanding such a text?

I argue that the problems associated with induction can be solved, or at least greatly clarified, by taking on an ignorance-first, outside-in, inverted view of epistemology. The riddle of induction is the legacy of a positive epistemology that conceives of inductive inference only as the cumulative aggregation of positive, experiential contents that upon reaching some mysterious critical tipping point of gathered case examples, extends or maps to the unobserved through an inductive generalization. The perspective changes when we view the relationship between knower and unknown as a kind of thermodynamic work cycle that is two parts of the same coin and one system. Knowledge, by these lights, is understood as a *far from equilibrium metastable order*. Whereas high ignorance is related to what is called the *principle of maximum entropy*, which states that uncertainty is highest when all outcomes are equally probable (Jaynes, E. T. (1957). "Information Theory and Statistical Mechanics" (PDF). Physical Review. Series II. 106 (4): 620–630. Bibcode:1957PhRv..106..620J. doi:10.1103/PhysRev.106.620. MR 0087305). The metastable order maintains its stability until it is perturbed. Learning occurs when the equiprobable distribution within the space of ignorance suddenly spikes with an irregular change (surprisal). Such a spike of the unexpected spurs a change to the inferential model. The inferential, generative model which continually tries to predict its inputs, is not destroyed by this unexpected spike, that is, the discovery of a black swan, it is instead *revised*. This view that an inductive inference, or rather the process that yields one, must be flawless to be valid, is an archaic and rigid conception that pays no respect to more flexible Bayesian understandings of the belief revision process. The truth of inductions is provisional, always pending further inputs.

When the unexpected spikes knowledge moves to fill that gap. knowledge must break so it can reform into a more robust framework. The agent is driven to epistemic action by the recognition of its own ignorance prompted by the surprise. The spike in the unexpected triggers the cognitive work cycle in which the inferential engine springs into action. It activates and drives toward the source of the spike until doubt is resolved and stability in belief is restored. Thus, in a process reminiscent of energy flow over a gradient, as knowledge

accumulates ignorance reduces. Learning is *perturbing*. As DeNicola writes, "learning may disrupt our cognitive comfort; it displaces us. Education requires us to revise or abandon our routines…and to do so we must overcome a kind of natural cognitive inertia." Yet it is precisely this metastable, steady-state of resting inertia which is the goal of all thinking, the Piercian fixation of belief. Doubt is an irritant, and we'd prefer not to have to deal with it. Thus doubt is a spur to action. Belief is perturbed so that it can resettle on a new stable state that has incorporated the new information.

A perfectly known world is one of stasis, a world that is effectively timeless and changeless. Cognitive activity might as well cease in such a circumstance of perfect stasis, and paradoxically, the performance of the knower would degrade over time because it would no longer encounter any challenges. Ignorance is the driver of knowledge, without it there would be no knowledge. Another way of putting this is that a perfectly known world must be a perfectly simple world, one that can be encapsulated in a few trivial observations. A world that simple however would not support sophisticated intelligence. Intelligence withers in the face of the simplistic. The situation is comparable to Rosenzweig and Bennett's famous study of neuroplasticity in lab rats placed in stimulating, enriched environments versus bland, impoverished environments. The rats in the enriched environments with a hamster wheel and other bells and whistles later exhibited bigger, more densely interconnected brains. Meanwhile the rats in the spare, spartan cages with barely anything to do showed the opposite. In the parlance of Hebbian learning on which brain plasticity is based, it comes down to a matter of "use it or lose it." And you only have to use it when encountering a complex, variegated world. Otherwise you wouldn't even lose it. You wouldn't have it in the first place.

Intelligence is a response to complexity, it is an adaptive attempt to *meet the challenge* of complexity. The more complex a thing is, the more ignorance there is initially about it, because the more it takes to describe it starting from scratch. Therefore the fact that induction is error prone is a demonstration of the challenges it faces, it is evidence of the interestingness of thee subject about which it tries to infer truths. Where it too easy to infer what is going to happen or

what is true of a class of things, it's probably not something worth thinking about.

Of course the complexity of an object is not always immediately apparent. Our own bodily functions are an example of this. To us walking is among the easiest things to do. Subjectively it's exceedingly simple even though if we were asked how to explain how it happens in detail, a common sense answer would be hopelessly lacking. "I put one foot in front of the other." On one level that answer is satisfying, but it hides the complexity at play. Objectively the detailed description of bipedal locomotion is excessively complex, involving many muscle groups in coordination with many bones to control fine motor adjustments and gross movements; it even counterintuitively involves vestibular structures inside the inner ear that regulate balance. And all the while this is controlled and mediated by a complicated array of brain centers, nerve fibers and nerve endings.

It is possible to see the complex as superficially simple and therefore never be prompted to learn. Unless one has a reason or desire to or is otherwise forced, things are often taken for what they seem at first glance. The uninformed novice can look at something and see only an undifferentiated mass whereas an expert sees rich details. Consider a skilled sports commentator naming all the technical terms of an olympic ice skater's maneuvers, whereas the average person is unable to discriminate such granular detail. All they see is various barely distinguishable leaps and spins.

Often this insensitivity to finer points of hidden complexity comes to an end when, according to Heidegger, we lose the *readiness-to-hand* of an object. Our bubble bursts. Placid obliviousness as to the true depth of an object is destroyed the moment the unconsciousness of it is interrupted by a discontinuity, breakage, removal, or change in the object and therefore our assumptions about it. Consciousness of an object's true being comes into play when it contradicts our inertia toward it. We are not conscious of the true details of walking because we don't have to be. It is not even something to think about. If one is injured and needs to be rehabilitated to walk again, suddenly we are more aware of what exactly goes into it. We have to re-experience what it is to learn how to walk.

Precisely such a break with locomotor normality might spur one individual to study it more deeply, until the physiology and biomechanics of it is understood and a science developed. And so it is with diseases, and all the other problems of life. They seem to be simpler on the surface until we are forced to think about them like the nail that sticks out and gets hammered. The same goes, for example, for the typical person using a software application. As the user we don't see what is "behind the buttons" even though a tremendous amount of technical operations and logic might unfold the second we hit effortlessly send.

The perspective shift on induction I am advocating for is one that puts ignorance before knowledge. Being wrong when it comes to induction does not mean that the entire model must be thrown out. Upon encountering a black swan the inference that all swans are white becomes, "all swans up to this point were white and there is at least one black swan." There is no good reason to rigidly commit to a static, unyielding idea of induction that must be true all at once. Inductive inferences are an ongoing process, just as the universe is. And in fact, the very idea of induction makes no sense without a temporally extended series of accumulated confirmations of past experience. Cognitive agents are constantly bumping up against their ignorance and updating their world model accordingly. Rationality updates and adjusts their model in response to the discovery of the previously unknown. Hume's problem of induction refuses to recognize this incrementalist view of induction. He assumes that the facts are what they are, that all facts are given at once. The uncertainty that induction deals with is not a flaw of induction but a property of the field of potentials it operates over. Unfortunately when it comes to guessing what this world has in store there is no choice but to gamble.

This all swings back to empiricism and its boundaries. Recall that the the position that all truths are derived from experience is not a truth derived from experience. The hypothetic-deductive method championed by Popper is intended as a response to this inconsistency. According to Popper the true model of scientific reasoning is a kind of inverse of the more conventionally embraced induction. Empirical activity begins with ideation, with a guess about the phenomenon in question, and then an experiment is developed to test the

hypothesis. The result then feeds back into the guess and either spurs the scientist to reject or preserve it. Popper tacitly agrees that uncertainty is irrevocable, by supporting a falsificationist view of the scientific process where the truth of scientific claims are not provable but pending falsification (Popper, 1963, Conjectures and Refutations). We cannot say that a scientific claim is proven, but we can say it has not yet been falsified and that the claim is worth accepting because it has survived an onslaught of attempts to falsify it. A scientific theory is like an unbeaten boxer, who might one day meet his match. We cannot rationally conclude that the boxer is unbeatable from their record of victories (non-defeats).

Combined with a pragmatist view of truth, scientific theories enjoy distinguished epistemic status for two reasons: they are durable in the face of efforts to disprove them, and they are useful. The have both negative and positive validations. They have remained standing through it all, and they solve problems. This is as good a reason to believe something as anything.

If one takes the hypothetic-deductive method and applies the same verification standards as induction, you will see that it is subject to the loopholes and the same climate of epistemic uncertainty as induction. The method of guessing and testing is not itself subject to guessing and testing. (Perhaps here the ghost of Gödel is lingering in the shadows.) The omnipresent specter of uncertainty-ignorance is just as present for this method which is proposed as a more accurate depiction of scientific inference. It begs the question, what exactly goes into the guess, how do we arrive at a good enough guess to arrive at a test which rises to the top of the heap of the not-yet falsified? Really what one is doing here is trading one set of conceptual deficits for another, because in reality: the flaw is not contained in the method but in the fundamental limits to the cognitive legibility of the world.

As Hume observed the concept of causality cannot be inferred other than by affirming the contiguity and similarity of different sets of events and conditions. “We may define a cause to be an object followed by another, and where all the objects, similar to the first, are followed by objects similar to the second. Or, in other words, where, if the first object had not been, the second never had

existed", he writes in the *Inqury*. Observation can only confirm that relations of similarity hold between contiguous objects. A causal model cannot be arise from these relations, only the inference that certain events tend to co-occur. Though an intuitive deductive method seems to be more powerful because it provides a more clear explanation for how theories are developed, its inner processes are not obvious or publicly scrutable. In either case, we are dealing with a "unknown known", unconscious tacit knowledge, an un-inpsectable hidden stow away inside the mechanism of inference. This is bad because when an inference is wrong there's little to tell us how the approach that led to the wrong conclusion differs meaningfully from the one that didn't, besides the outcome. James, in line with this thinking, points out that a belief is not true by itself, it awaits confirmation from the world. "The truth of an idea is not a stagnant property inherent in it. Truth happens to an idea. It becomes true, is made true by events. Its verity is in fact an event, a process: the process namely of its verifying itself, its veri-fication. Its validity is the process of its valid-ation." (James, *Pragmatism*, 1905). Such a notion of truth happening to an idea rather than the idea being true in the first analysis is in line with ignorance-first ontology and epistemology. It is quite hard indeed to be certain about something that from our perspective doesn't yet exist. "..we have only one edition of the universe, unfinished, growing in all sorts of places, especially in the places where thinking beings are at work." That the universe is "growing in places" where cognitive agents are at work, is an view of cognition wherein the agent is actively participating in shaping its environment such that certainty is more possible for it. Where the universe is growing, in James' turn of phrase, information entropy is being reduced, possibilities are collapsing into definite states. Indeed the universe *is* in fact literally growing. As current scientific consensus holds, an abundance of evidence shows that the scale of space is inflating such that the relative distances between cosmic bodies has grown in time and is indeed accelerating. This kind of growth is not quite what James meant by it, but it helps illustrate the idea that the universe's possibilities are not yet fully actualized but are unwinding like a rolled up carpet.

This discussion about inference may come across as belabored but that's

because it must be. On one view inductive inference is a fundamental part of the human condition and experience. It is no abstract philosophical curio, but the central feature of the mind in relation to the world. Induction is how experiences are knitted into a coherent whole. It's the spear tip that pierces the unknown. It would be a garish, overgeneralizing inductive leap, however, to assume that we can use it to explain everything. That's why this analysis brings us naturally to the next topic.

The deep connections between uncertainty and possibility go beyond inductive inference. They are embedded in functions of the brain and in the developmental cycle of human beings—and more generally, in evolution and life at large. Significant insights into the study of ignorance can be had by examining the structure of the brain, its growth patterns, and how its plasticity and adaptability fits into the picture.

Part IV : Ignorance In A State of Nature

Our wide ranging exploration of ignorance has brought us into contact with many different perspectives. When all is said and done however, the irreducible fact remains that ignorance is most strongly associated with minds and brains. At the same time the process which produces life, evolution, itself seems to embody ignorance-theoretic principles and this is no contradiction because again, ignorance is in some concrete sense objectively mind-independent even if minds are what interface with it most directly. There is also the question of naturalizing ignorance, that is, opening it up towards the sciences of psychology and so forth. This whole work in large part is a project of naturalized epistemology, or rather agniology, but it is helpful to explicitly articulate why that matters. Ignorance is no ethereal matter, but a fact of life just as decisive and impactful as any. It helps then to treat it as something which can be examined as precisely as any scientific variable.

Agnoiology Naturalized

In pursuit of a more comprehensive theory of epistemology that includes a theory of ignorance and inquiry, many threads must be investigated. It is not enough to stop at the abstract contents and relations of epistemic states considered a priori. We need to put some empirical flesh on these theoretical bones. Nor can we be satisfied with an additional coat of paint detailing the physical thermodynamic and informational correlates of these logical abstractions. We must have a thorough understanding of the engine of uncertainty-reduction as a concrete empirical object, the brain.

The idea of naturalizing epistemology is not a new one. Clear traces of it can be found in the empiricism of Bacon (both Rodger and Francis), Locke, and Hume and its development can be followed in the gradual transformation of that empiricist philosophical program into the underpinnings of modern scientific psychology. In general, naturalistic epistemology states that epistemic concepts are best analyzed as methodologically, substantively or otherwise explanatorily tied to scientific approaches, theories and findings. Epistemology, rather than inventing its own bespoke terminology and conceptual framework, ought to recruit and incorporate science and its explanatory materials and be expressed in those terms. This naturalistic paradigm is certainly implicit in the methodology used in this book, but now comes the time to take it to the next level.

A touchstone for the naturalistic approach to epistemology was Quine's 1969 paper *Epistemology Naturalized* in which Quine explicitly laid out the argument. Initially Quine supports the idea of rational reconstruction in the vein of Carnap's *The Logical Structure of the World*, in which sensory experience is to be translated into well-defined *observation statements* which are then subject to logical operators, allowing one to schematically construct a logico-mathematical "language of science." While yearning for such a scheme, Quine accepts that is untenable. In the most basic terms, such a translation procedure would require that "blocks of theory" or theories as a whole would have to be translated into observation statements, but this is impossible because theories are not direct transcriptions of any one empirical observation. There are parts of the

apparatus of a theory that do not correspond to any particular sense input, and therefore an asymmetry holds between theory and observation . Nor would a case by case collection of observation statements for each empirical data point constitute a theory. Rather, this is just a data set. The theory consists of a framework logical and causal relationships that are superimposed on the data. If there is no away to logically reduce sensory evidence to use it as a logical model for epistemology, what basis do we have for an a priori theory of knowledge at all? Why then, Quine asks, do we not "settle for psychology"?

Yet there are many reasons to resist Quine's proposal that epistemology be eliminated, or "outsourced" to psychology, or be considered as a mere chapter of its book in the library of science. As Reichenbach said, "epistemology does not regard the processes of thinking in their actual occurrence; this task is left entirely to psychology." Rather, epistemology seeks to identify what makes certain beliefs or mental states consistent, coherent, or otherwise hold distinct logical properties such that they are distinguished from other beliefs by their special status as knowledge. (Reichenbach, Hans, 1938, Experience and Prediction, Chicago: University of Chicago Press). This is not the place to review these critiques exhaustively. The goal is only to summarize the history of the naturalistic epistemological project. Naturalized epistemology in a broader, richer sense than Quine's eliminationist stance, states that the distinguished logical properties that belong to knowledge must either supervene on natural properties or are best analyzed in terms of such phenomena. Epistemology, regardless of any formal considerations, has much to learn from the sciences.

The next step in a naturalized epistemology is to pick which sciences are the most relevant for descriptive and analytical purposes. The most obvious is psychology in its various iterations. Information theory and its associated related concepts in computer science and physics are resources which have already found ample usage throughout the present analysis. Artificial intelligence, which has also found some traction, is another good choice. Each is a worthwhile addition to the toolkit. Among the most compelling options however is neuroscience.

The progress of neuroscience has made extraordinary leaps and bounds in

recent decades. It straddles the middle ground between purely empirical studies of functional brain structures and the underlying mental states they cause. Among the pioneers of neuro-philosophy is Patricia Churchland. Churchland, in a variation on Quine, suggests that traditional epistemology is contaminated by folk psychological, that is, everyday, scientifically "untutored" common sense explanatory assumptions about mental causes, events, states, or actions (Churchland, 1987, Epistemology in the Age of Neuroscience). This muddle of indistinct folk concepts ought to be cleared up by the description of a more sound conceptual system framed by empirical findings and the language of science. Yet is is not immediately clear how we can go from a three pound glob of brain matter, or from microscopic, spindly neurons, to a comprehensive understanding of what knowledge is. Churchland and Churchland (1983) broadly differentiate between two basic approaches to this quandary. There is a computationalist approach, which essentially views the brain as a symbol processor and emphasizes the "linguistic and rule-following" aspects of cognition. Then there is an evolutionary, neurophysiological approach, favored by the likes of Darwin, Helmholtz, and Hebb which stresses biological drives, instincts, functional brain areas, genetic predispositions, innate cognitive modules, adaptations, and a similar repetoire of explanatory variables. Hebbian learning, for instance, describes a principle of neural activity in which repeated stimulation of a post-synaptic cell by a pre-synaptic cell cause them to establish a persistent link ("neurons that fire together wire together") in a process that physiologically underserves the psychology of learning. These computational or evolutionary approaches are not by any means inconsistent with one another, but rather specify two different naturalistic research programs that intersect often.

For the present purposes the obvious point of departure is to assume that if the project of naturalistic epistemology is a fruitful enterprise then a naturalistic agniology must also. If in other words the thesis that uncertainty/ignorance is an epistemological ground condition is true, the brain ought to be organized to reflect its teleological goal as the engine of information entropy reduction. This organization is represented not only in what the brain *does*, but in how the brain *is*. Structure is function. Moreover we should be able to follow this line of

reasoning even beyond the brain into other natural systems.

This project is a considerable one, most likely deserving a volume of its own. The objective here is to provide, in reasonable detail, the outlines of such a project.

From the outset it should be noted that the analysis sketched here has a much more comprehensive view of its subject than than mental states and other psychological factors. The objective is to really lean into the naturalistic standpoint and to draw from an expansive perspective. Uncertainty is best thought a pervasive force and as a natural kind, one that goes beyond mental states and is embedded in fundamental physics, as such, the investigation must be framed beyond the traditional boundaries of epistemology. In consequence entities which are not conventionally seen as epistemic concerns will be treated as such.

Three areas of principal interest for a naturalistic study of ignorance, the neurobiology of learning and development, evolution by natural selection, and immunology, can be delineated. While evolution is the most fundamental in the sense that it caused the others, the brain's mechanisms are the most immediately relevant, so it will be addressed first. Immunology, while an interesting topic in its own right, is posited as an intermediate level between the other two that nevertheless embodies the same principles of learning in the face of uncertainty.

More can happen than does happen. At least that is the generally accepted premise of thermodynamic physics with the concept of *phase space*, which denotes all possible values of the position and momentum variables within a particle system of which at any given moment only a single one of its states is manifested. In any finite span of time, only a tiny subset of all possible state configurations of a gas will be realized. This is a nifty way to think about macro-scale objects. Admittedly it's not exactly an established position within physics that what is applicable in the micro must also apply to the macro. For example while there is the concept of the Heisenberg cut, the proposed theoretical spatiotemporal boundary at which quantum effects are supposed to yield to classical effects, this boundary is merely theorized rather than observationally

confirmed. And there are examples of macro-scale phenomena, superconductivity, superfluity, and Bose-Einstein condensates, which demonstrate that this boundary is permeable. In any case, statistical mechanics is described using classical physics, so if the particles are being treated like billiard balls, the billiard balls can be treated like particles.

Thermodynamics after all encapsulates the physical basis for all change and motion. For anything to happen in a brain, including presumably the beliefs represented within it the that constitute knowledge, energy must be expended. Because anything could happen in the phase space of possibilities, any theory of a cognitive agent must be account for this principle in terms of how the brain structures itself in response to novelty and the unexpected. You cannot be rigidly adapted to just one future outcome in a complex world that is best thought of as suspended in a state of ambient possible outcomes; the very act of rationally weighing decisions presumes the awareness that different outcomes are conditional on different causes of action. The brain, therefore, comes inbuilt with a plastic, flexible connective structure specifically to adapt to come what may rather than to strictly assume what will come. The brain on a fundamental structural level is evolved in view of expected uncertainty.

A dual nature is exhibited in the brain. Not only does its function reflect uncertainty reduction, so too does its structure. This flexibility is most evident in the grow-and-prune mechanisms of neural development.

The neonate starts with a brain that is highly modular in its overall organization and the general spatial allocation of neural populations, yet flexibile and protean in fine detail. Neural resources are overgenerated from birth. Newborns have twice as many brain cells at half the overall brain size of adults, and acquire them at a rate of 250,000 new cells for each day of pregnancy. A point-for-point allocation of each cell precisely to where it must it go and will end up being used in adulthood is a level of precision which for such a complex system perhaps exceeds the boundaries of biological possibility. The amount of information needed to specify that finished state is much greater than what the genes actually specify, the general rules for brain development rather than the outcome. The developmental sequence is one of "great orderliness and efficacy

but also allow\[s] room for the creative effects of chance." As such in the developmental neuroscience literature it is standard to propose, "a series of stages in which built-in instructions and the effects of arbitrary external events are mingled to an intriguing degree." (Ackerman, Wes, 1992). The developmental stages are as follows:

(1) proliferation of a vast number of undifferentiated brain cells; (2) migration of the cells toward a predetermined location in the brain and the beginning of their differentiation into the specific type of cell appropriate to that location; (3) aggregation of similar types of cells into distinct regions; (4) formation of innumerable connections among neurons, both within and across regions; and (5) competition among these connections, which results in the selective elimination of many and the stabilization of the 100 trillion or so that remain. These events do not occur in rigid sequence but overlap in time, from about 5 weeks after conception onward. After about 18 months of age, no more neurons are added, and the aggregation of cell types into distinct regions is roughly complete. But the pruning of excess connections—clearly a process of great importance for the shape of the mature brain—continues for years. (Ibid)

Learning related synaptic pruning happens in the maturing brain and continues until around the mid-twenties, during which the initially overstuffed neonatal brain loses half its cells but grows twice in size. Terminal arbors (the end points of the projecting cell) are selectively pruned through differential use and disuse that recalls darwinian natural selection (Edelman, 1987). This pruning effectively "sculpts" the initially undifferentiated microstructure of the brain into a more functionally coherent architecture based on the shaping effects of sensory stimulation and local activation patterns.

Why does the developing brain follow this plan? One would think that babies, with all those extra cells, would be smarter than adults and we'd only get stupider in time as half our original brain density withers away. Especially given the encephalization quotient: brain size tends to be proportionate to body size. Newborn babies' brains is half the size of an adult's, but their bodies are not; rather, their brain is bigger than their body size should predict, and the size of their crania are much larger proportionate to their body size. This is sensible

given one important developmental detail: compared to every other organ the brain is the most incomplete at birth; it's the only organ whose incompleteness does not pose a serious congenital defect. It's a feature rather than a bug. In some sense it is true that babies are smarter than adults: children learn language and other complex tasks with an ease that defies adults. Their brains do not engage in the higher level skill acquisition that adults do, but they are learning the more primary tasks of how to see, hear, orient the body, walk, categorize perceptual objects, and understand language. Each of these is quite the tall order, and so it makes sense that they would be provisioned with ample neural resources. Infants accomplish the monumental feat of putting together a world.

A grow-and-prune model of brain development is exactly what one would expect if uncertainty were the motive principle behind all cognition and that trial and error the most primordial means for coping with it . The under-provision of neural resources is obviously not the right option. You would never produce fewer brain cells than you would need to make a functioning human being. Neither is supplying "just enough" cells, as this provides no wiggle room in a world where anything could happen over the course of growing up, where mistakes will be made and a momentary lapse in oxygen can kill off droves of neurons. When building a house you do not show up with precisely as many nails, wooden planks and roof tiles as you would eventually need. The supply comes padded. You bring reasonably more than you would need to account for accidents, mistakes, spoilage, miscalculations, and last minute changes to the blueprint. In effect, you must cast the net wider than whatever it will close around.

Over and above the logistical problem of providing some wiggle room for the construction of the maturing brain, there is another problem of the descriptive complexity, also known as the algorithmic entropy or Kolmogorov complexity, requisite to express the structure of a mature adult brain (See Ruffini 2007 for a good overview of the descriptive complexity of the brain). The descriptive complexity of an object is the minimum length of a program needed to generate the object. To use a computing analogy, the genome is like a compressed zip file that when unzipped expresses the organism to which it pairs. Put in the simplest

terms, the sum of information the adult brain will end up needing to operate effectively is greater than the amount that the genome can initially encode. If it weren't so, it would be reasonable to assume that either children would be born with the complete cognitive functioning of adults, or mature much faster than the unusually long developmental period typical of human beings. But this would mean humans wouldn't be able to do as many interesting things. It takes time to build up the missing portion of information needed to live an adult human life that's left out by the genome. Human brains spend about a third of the lifecycle maturing. Mice in contrast live about two years and mature in six weeks, meaning that they spend just 6% of their lifespan maturing. One can conjecture the reason for this developmental brevity is that the descriptive complexity requisite to express the simpler mice brain is much less than that of a human's.

The human genome "knows" that it cannot know what the maturing organism will encounter, such as what specific language it will need to learn, and so exactly specifying the configuration of neural resources could bring the system into contradiction with its unpredictable future. To exactly specify the future would have to be an exact copy of the past. The multidimensional bitmap representing the mature brain results a fractal pattern that cannot be stored in the lower dimensional compressed state of the genome. The genome can only provide the parameters that get the ball rolling, it sets the rules but does not play the game. It specifies, in other words, the abstract class rather than the concrete instantiation, Chomsky's instinctive universal grammar rather than the grammar of a given language.

The difference between what the genetic endowment grants as a given—reflexes and instincts, general rules—versus what is learned freely, is an indirect statement about the degree of predictability of the world. The withdrawal reflex, sweating, shivering, and sneezing, or the various primitive (infantile) reflexes like the Moro response, are uniformly inherited fixed action patterns which constitute a kind of evolutionary "bet" that a certain set of situations will be encountered in the course of a lifetime. Namely situations where a limb must be rapidly retracted, or for when it is too hot or cold to maintain internal body

temperature, or for when a newborn is startled and must cling to its mother. The cards come stacked with these givens as there is no probable future where these responses don't apply. Animals with free, cerebrally meditated learning capabilities, however, represent a peculiar advance over this reflexive toolkit. It enables proactive knowledge-seeking. Humans have domain general cortical learning. They can actively detect features of their environment which their genes could not account for and work around them by devising novel behavioral innovations. The so called human *cognitive niche* which "embraces the ability to overcome the evolutionary fixed defenses of plants and animals by applications of reasoning," emerged as a kind of dialectical overcoming of precisely those fixed patterns which generally confine the behavior of other species (Pinker, 2010). This effort to solve the various puzzles posed by other species' ingrained defense mechanisms drove the "outside the box" thinking of human cognition which was no longer bound to strict situation-action conditional reflex behaviors or instinctual programs. Given that human culture and other people also comprise the human niche, this selfsame unbounded cognition drove itself to further heights as human ancestors strove to then outthink and outwit each other at their own game in an escalating arms race of competitive cleverness.

Taken together these details suggest the brain's selectionist design plan is attuned to an environment of expected uncertainty (for a fuller and originative account of neural selectionism, on which this account is based, see Edelman, 1987). Corresponding to its open-ended, connective potential stands the open-ended possibilities of the world. This physiological process of synaptic growth and pruning and the successive adjustment of neural resources to a target environment is adaptation to uncertainty. In a manner of speaking, it is the *adaptation of adaptiveness*, second order adaptation. Rather than a fixed, adapted feature like a elongated bird's beak to fit the peculiarly cylindrical petal shape of the corolla of a flower in a particular island habitat, the adaptive learnability of the CNS is *itself* an adaptation. In the same way that the weather is to the climate, so individual learning is to natural selection. Both learning and evolution are adaptive responses to short term and long term change,

respectively.

In development, neural overproduction is met with selective elimination ("use it or lose it"), a multistage process in which successive waves of growth and pruning expand and then refine structures until long-term functional circuits are established. This meeting in the middle of continually readjusted internal functions to external inputs assumes from the get-go that there is some level of uncertain variance in the input. Such an architecture posits an openness to a phase-like "state space" of branching paths each corresponding to the various environmental contingencies the organism may encounter. A given ontogenesis therefore contains information about a given experiential history, as everything the organism has done leaves a trace and biases what it will do in the future. In rodent studies, mice that had artificially boosted neurogenesis (brain cell production) exhibited increased perceptual skills over controls, suggesting that having "more to lose" is beneficial to the learning process (Fang and Yuste, 2017).

Every act of inclusion is an act of exclusion. Much like the negative definition of knowledge acquisition described earlier, neural selectionism must destroy to create. The mature brain is a residual reduction of everything it is not, it's the exclusion of alternate path histories. Its concrete structure is the product of shedding the possible for the sake of the actual. The overproduced neural populations of the newborn cast the widest possible net, accounting for a phase space of alternative life histories only a subset of which will come to pass and end up encoded in the finalized, pared down neural network. Accordingly, what we can do is dependent on what we have done, which becomes a self-reinforcing loop. Counterintuitively, therefore, the ability to do more coincides with an overall decrease and loss of potential. A more defined behavioral repertoire comes into focus through the exclusion of a less defined, yet more potent space of possibilities. Greater organization only comes to pass through a more abundant though less organized initial state. The neural ontogenesis recalls the concept that a sufficiently rich randomness contains useful ordered subsequences and the difficulty is the cost in detecting and extracting the patterns. Brain networks are far from randomly distributed however. It is marked by a careful cadence of developmental timings, in which, "neural stem cells undergo progressive

restriction in their competence to generate different types of progeny in response to extrinsic signals." (Xin Li, ... Claude Desplan, in Current Topics in Developmental Biology, 2013). Such nonrandom organization takes shape, however, within a higher-order state space which, for all its complexity, could be considered effectively random. Yet this developmental timing process and the nonrandom structures it produces is more akin to a controlled chaos, a stochastic mechanism in which noise imparts freshness to the system. This noise that breaks deadlocks and softens overly rigid patterns. Noisy brain activity, "breaks deterministic computations," and such symmetry breaking, as with the indecision of choosing between two equally valid options, can prove pivotal in decision making. (Rolls and Deco, 2010). Making "a decision that is not optimal based on previous history may provide information that is useful, and which may contribute to learning." An organism, for example, which is strictly deterministic in its behavior will not randomly sample its environment to test if is changing. It will lose flexibility. A degree of randomness, therefore, saves the cognitive system from freezing up.

The unifying thread throughout this account of the brain's developmental patterns and organization is that it embodies its own expectation of uncertainty. It (tacitly) knows it does not know what its finalized mature state will be, and accommodates for this structural ignorance through an open ended design rather than by the rigid imposition of a overdetermined set of growth pathways that correspond to only one set of anticipated environmental contingencies. The way of the world does not allow for such a brain. The genetically determined open plan of brain development is not like one might expect to see in other, admittedly less complex organ systems, such as the development of the skeletal anatomy, in which any variation from a strict blueprint is likely harmful. To become the engine of knowledge, the brain must first be an engine of ignorance. It must be prepared to respond to changes that cannot be predicted in advance.

Given that the brain is the most obvious mediator in the relationship between knowledge and ignorance, it's intuitive to consider it the most central feature of a naturalistic study of ignorance. Yet as the title of Edelman's *Neural Darwinism* implies, the plasticity of brain development owes itself to the more general

principles of biological evolution. The strongest version of the argument for uncertainty is that it's not merely a property of a confused mental state, but is "baked into" the fabric of physical reality. Accordingly, any complex ordered structure that persists in time, and which gains order and complexity, must ultimately come into contact with the aperiodic irregularities of its environment.

Just as the brain embodies a model of ignorance through its developmental patterns, there is also the other side of it. Any system that emerges from disorder must eventually return to it, and so to understand ignorance naturalistically we must also understand senescence.

Not surprisingly children are simultaneously the most ignorant and yet the most rapid learners. Growth naturally equates to learning. All else being equal— assuming the cause of one's ignorance was that one was just born or one is otherwise new to something— a state of maximum ignorance is also a state of maximum learning potential. In this condition everything one learns comes as a freebie, you simply have to exist in time and you will come to learn something new. As one learns more and ages, however, the learning rate starts to slow, with all else considered for no reason other than the laws of probability and notwithstanding the fact that the brain experiences decrements in cognitive function as it ages. Early on the odds of encountering something new and thereby learning from it is highest, but over time one is bound to come across informational redundancies and thus lower odds of experiencing novelty. At a certain point you've heard it all before. More than just probability is in play, however. You of course haven't literally heard it all before, it's only that the rate of neuroplasticity has slowed. The neurobiology of learning changes as one ages (Erickson, Barnes, 2003). Generally cognitive flexibility, the faculty of thinking on your feet and quickly switching between concepts decreases with aging, while crystalized knowledge—accumulated expertise—increases. The white hot cast that was the early brain cools and begins to harden into a solidified ingot. Eventually, tragically, there will come a point where senescence takes its toll and the quicksand of ignorance from which we spent a lifetime emerging out from claims us in the end as cognitive faculties decline. A pedant might split hairs

here: is cognitive decline really a cause of ignorance, or should there be an exception made for the physiological disruptions or dysfunctions that induce forgetfulness, misremembering, and other errors? With the addition of any cognitive error that is not recognized as such and corrected, there is always a corresponding increase in ignorance. If the medium on which knowledge is stored becomes corrupted, so does the knowledge with it. Therefore age-related cognitive decline is what one might call *agnoiogenic*, ignorance inducing. (More will be said about this concept of ignorance production in a later chapter.)

Of great interest is ignorance in the dimension of time given that time defines the direction of information flow and ignorance is all about issues related to information access. Previously we spoke of ignorance in relation to the physical aspect of time, and how there are informational asymmetries between past, present and future, but more can be said about the subjective sense of time and how it impacts how the mind works. When we speak of someone as "stuck in the past" we are implying that they are a certain kind of ignorant, that their mental model is out of sync and out of date with the current way of the world. Their minds, in other words, cease to self-revolutionize in the face of ongoing change, fixating instead on locked-in representations keyed to a bygone era that existed when their brain was initially developing and actively encoding a model of the world. This "cognitive fossilization" is the end result of cognitive sedimentation. As Nietzche put it, "the snake which cannot cast its skin has to die. As well the minds which are prevented from changing their opinions; they cease to be mind." Eventually all our built up knowledge, rather than simply negating ignorance as one would expect it to, becomes the one hammer we have which sees everything as a nail: we can't help but see the world in one stiff way, even if it admits of many more valid perspectives of which we are thus fated to be ignorant. Outdated knowledge eventually closes around us like a hard shell, as we become habituated to a certain kind of thinking at the exclusion of fresh alternatives. And while amenable to some degree of mitigation through lifelong education and proactive learning, this self-limitation is a fate which befalls all. To keep up with things requires a constant state of updating one's beliefs and concepts, a task which is easier to do when young and there are no

preconceptions but becomes increasingly difficult as the rate of change slows in the mature brain and its wiring in broad terms has already been determined. Eventually the intervals between which beliefs update become wider and wider, and less impactful, until eventually they cease updating nearly altogether.

So the curse of youth with its sterotypical naivety and inexperience and brashness (all expressions of ignorance in their own right), eventually gives way to the curse of age: failure and slowness to adjust to changing circumstances and a stubborn bent toward mental conservativeness. If the goal is a naturalistic view of knowledge and ignorance, then an idealized abstract knower that floats comfortably above all these mortal frailties can't be assumed. Much can be gained from what Lakoff and Johnson dub *philosophy in the flesh*, the embodied underpinnings of the mind that prefigure reasoning. knowledge, and by extension ignorance, are shaped by physiology and bodily experience. They do not occur in a transcendental realm of pure reason detached from the state of the body. How one's relationship to ignorance changes over the course of a lifetime is therefore an important piece of the puzzle. Like a cataract, senescence clouds over the mind with false certainty. It creates a subjective frame of reference around which everything must either fit or be rejected regardless of the real truth in play. Vexingly, the more certain and set in our ways we become in time, the more a devious form of ignorance and insensitivity to change risks creeping over us. You can't teach an old dog new tricks. Indeed, this conclusion is backed up by the science. There is an inverse relationship between a subject's age and the amount of effort to change neural circuits in response to new experiences: younger the easer and the older the harder it is (Levitt and Eagleson, 2009).

Perhaps from a generational standpoint such age-related dynamics help explain Kuhnian scientific revolutions. As physicist Max Planck in a bit of dark humor once quipped, "science progresses one funeral at a time." According to Kuhn, scientific theories follow a cyclical pattern of acceptance, rejection and replacement. A specific theory comes into prominence, and an old guard consisting of those researchers who staked their career on that paradigm defend it by resisting the counter-evidence against it or explaining it away. Meanwhile a

hungry, younger generation of scientists see the holes in the reigning paradigm of "normal science" and seek to make their name by overturning it. As a research program continues, the number of anomalies the outdated scientific consensus fails to address becomes unsustainable. The newer generation of scientists storm the ramparts and spur a paradigm shift that establishes a new theory that accounts for the anomalies and on it goes.

It might indeed be a fruitful exercise to interpret Kuhn's *The Structure of Scientific Revolutions* in light of the theory of ignorance. Kuhn's proposed lifecycle of the scientific process could be reevaluated in terms of ideas such as structural ignorance. It could be that what happens during such scientific revolutions is the generational churn and turnover of various stages of cognitive spryness versus rigidity, of learning-facing the ignorance of outdated understanding versus outdated understanding facing its obsolescence. The young are naturally inclined to view things differently while the old see the value in the established and reliable. This contrast creates a push-pull tension that eventually results in scientific change in particular and generational change at large.

The so called curse of expertise is a variant of this age-related knowledge and ignorance dynamic. With age comes experience, but with experience comes an illusion of competence (Fisher and Keil, 2015). Expertise can form such an epistemic cataract,

The discrepancy between expectation and reality which carries the surprisal value—the difference that makes a difference which triggers a learning episode—prompts two different possible responses. Either the new information is accepted and integrated into a revised perspective. Or there must be something *wrong* with this new information, because how can I be mistaken? So it is instead ignored. It must be rejected or explained away to preserve the integrity of the cognitive status quo. In light of the expanded definition of ignorance this is indeed a type of ignorance even though it does not exactly signify a lack of knowledge. It signifies provisional knowledge which rejects other knowledge in contradiction with it. This denialism is actually knowledge— or what McGoey calls "strategic ignorance": knowledge that additional new knowledge would

break the spell of conviction, burst a bubble that would better be left unpopped. It is precisely this type of tainted knowledge which also goes by the name of *willful ignorance* or in James's parlance, the will to believe. The conceptual form of willful ignorance isn't defined by a lack insight about information contrary to one's position—in fact one can have complete access to this information—but the refusal to acknowledge that information. It's the awareness that the acidic touch of that information would corrode the edifice of held belief. As Freud was well aware, *self-delusion* is motivated reasoning that suppresses truth to preserve some psychological need or attachment. The result is the filter-funnel effect of confirmation bias where information that reinforces preconceptions is admitted while contradicting information falls by the wayside.

This discussion of willful ignorance will resurface later on where it will be handled in greater detail. Here it segues into an important research area of naturalized ignorance: cognitive bias. Cognitive biases are particularly interesting phenomena because they are specific ignorance inducing tendencies that reference neither specific mental content or the lack of content. They are not about beliefs, but belief formation. They are processes, leanings, that distort how information is processed as opposed to the consequences of a false belief or the absence of true belief. It takes a warped lens to put things in a false light.

We tend to view thinking and learning to be the straightforward routes to knowledge. The brain is a muscle, the saying goes, and the more you use it the stronger it gets, suggesting an uncomplicated, direct and positive relationship between cognitive activity and the growth of knowledge. knowledge acquisition is characterized by perfect lucidity and simplicity, it is a matter mere uptake as fluid and natural as a fish's gills that extracts oxygen from water. What if, quite to the contrary, thinking and learning were not such a naively pure processes, but instead deformed and warped by systematic distortions? What if the output is not anything like the input, and instead of being put through a linear assembly line, goes through a tangled and bending cluster of tubes where every sort additive and adulteration is injected into it on its way to becoming a settled part of the mind? And what if some input, no less valid than the others, has

something feature that causes it to be rejected along the way? Cartoon caricature though it may be, this latter description of the thinking process as a tortuous and muddy affair is more accurate. Enter cognitive bias.

Despite the identification of nearly 200 distinct cognitive biases and the efflorescence of literature on the topic, there is surprisingly little in the way of theorizing about what unifies them. What explains cognitive biases? Regardless of the exact nature of cognitive biases they are a most intriguing class of specimen because they represent the quintessence of structural, positive ignorance. For one, they are not deliberate or motivated self-deceptions like a selective memory or Freudian rationalization. We find ourselves committing them unwillingly and unconsciously, and find it hard to stop. Even when we learn about them and train ourselves to spot them, they still recur with the persistency of an ingrained bad habit. They are not mere factual misunderstandings that are rectifiable through education. Pervasive distortions of thought and judgement, cognitive biases are not caused by a simple lack of information or a single contingent mistake of reasoning, but are instead persistent side-effects or defects of the thinking process itself, kinks in the information processing pipeline.

Perusing the list of the dozens of cognitive biases identified by psychologists is uncanny. The more one learns about them the more one notices them at work in your thinking and-most often-the thinking of others, an observation which is itself a selection bias and a "bias blindspot" which is the bias that you are less biased than others. It makes you wonder just what aspects of the whole intricate contraption of human thought is spared from this malign influence. Since the discovery of cognitive bias by psychologists Amos Tversky and Daniel Kahneman, it has poked yet more holes in the Enlightenment rationalist picture of reason as the "mirror of nature" that can access the objective truth without interference or complication. Research has instead discovered that people's normal course of thinking is riddled with representational distortions. The key question is: is there a system to these errors?

The paucity of explanations for cognitive bias is partly accounted for by the lack of any universally accepted theory of unbiased, or *normative cognition*.

How can you have one without reference to the other? Also, a general preoccupation with how we are right, inherited from a rationalist, positivistic intellectual history, overrides any focus on how we are wrong, even though understanding both is required for everything to balance out. Perhaps, the trouble arises because the sharp distinction between normative versus biased cognition is faulty and has to be reevaluated. What if bias is a feature, not a bug? Is bias by design?

According to the argumentative theory of reasoning, the function of rationality is to win arguments. "When reasoners produce arguments, they are biased and lazy, as can be expected if reasoning is a mechanism that aims at convincing others in interactive contexts. By contrast, reasoners are more objective and demanding when they evaluate arguments provided by others arguments." (Mercer, 2016). The goal is not get at the truth for truth's sake but to convince other people of your claim. The unvarnished truth does not pick sides or have an agenda, an inconvenient fact for a species that is all about picking sides and having agendas. If this account of reasoning is true, arguers are incentivized to retrofit the truth so that it accommodates their drive to convince others. The social value gained from winning arguments and coming off as convincing outmeasures the value of being committed to the truth to a fault and even if it means undermining your own interests. Mercer points out that biases only appear as logical errors from the perspective of a logicist framework of reasoning that depicts it as a method of probability or expected utility estimation. In Kahneman's original work this model of reasoning was assumed. If The original motive behind reasoning is not to estimate objective likelihoods but to convince others of your position and thereby thrive in the social world, then biases start to look more like features instead of bugs.

A prominent example of such an argumentatively motivated bias happens to be one of the most well known and publicized: confirmation bias. Confirmation bias is the tendency to selectively seek out and interpret information that supports one's preconceptions. A logically equivalent definition of confirmation bias would then be to ignore all information that does not support one's priors. From a perspective on reason that views it as tasked with perceiving the whole

truth and nothing but the truth confirmation bias is deeply erroneous. Denial, the proverbial ostrich burying its head in the sand, does not negate the reality of what is denied. No valid verification mechanism counts ignoring the counter-evidence against a proposition as a method. If a model of reasoning is patterned not on the template of implementing abstract logical formulas of probability estimation, however, but on the practical considerations of winning a court case, suddenly the use value of confirmation bias becomes clear. No prosecutor or defendant in their right mind, who is trying to win a case that's stacked against them, gives the whole truth and nothing but the truth. The court system—deserving of study as an epistemic instrument in its own right—approaches the truth from a deliberate place of ambiguity: despite the logically necessary fact that only one side can be right, both sides have equal chance to make a case. Short of committing perjury a well presented legal case—especially when when representing a guilty defendant— is one that weaves a story around strategic omissions. Leaving out contradictory information simplifies the task of ramifying the connections between stated assertions. Every truth that goes counter to the argument is a workaround; the argument cannot include statements that would otherwise falsify the narrative, and so instead the evidence is presented in such a way that renders and account of events that differs from the facts but which still fits them well enough not to trip any alarms. The task becomes to constructing a believable argument that weaves around around the omissions. Legal trials are a paradigmatic example of argumentative reasoning because the goal is not to tell the truth but to tell the most convincing story. If the story happens not to fit the truth, so much the worse for the truth. The task is to make the truth fit the story.

Selective information gathering strategies are good ways to build a narrative that does not explode itself by taking on too much truth. An unflinching observance of logic and unswerving commitment to facts will not help in gaining social advantage at odds greater than chance. Only occasionally will objective facts shake out in a way that favors an unbiased argument. Despite the tiny fraction of it we have control over, the sum total of causality and the information flows stemming from it unfolds in complete indifference to human agency. Objectivity, by definition, is unbiased. A biased argument by its very nature has

better odds of being useful because it is a curated and cherry picked assortment of talking points. When it comes to convincing others the mind chooses to play with loaded dice rather than gamble with the facts. From this angle irrationality starts to resemble rationality.

The whole truth and nothing but the truth is disorienting. Truth writ large is not anthropomorphic, it does align with human needs. “People can’t stand too much reality” said psychologist Carl Jung. Inconvenient truths get in the way, pure objectivity, in most cases, is useless. Rarely does everything line up in such a way that every detail of our story bears out in pristine correspondence with the objective state of affairs. It has been said that history is written by the victors which implies a selective narration that edits and snips out various unflattering subplots.

As an intriguing aside, there is a variation on the confirmation bias driven by digital tech is the so called filter bubble. By analyzing past usage data, recommendation engines and search algorithms compute what they expect the user wants to see. But because showing them more of what they’ve seen in the past means a narrowing of what they will be exposed to in the future, eventually the algorithm converges on a sterile monotony. The net result in a long term impoverishment of opinion and a lack of novelty. Gone is the fruitful noise that shakes things up and brings a dash of the unexpected to the information stream which is what prompts a fresh bout of thinking. Diversity of signals within the information stream means more complexity, and with greater complexity comes greater intellectual challenge, which means richer representations and a more rounded, less watered down worldview. Filter bubbles technologically implement the cognitive tendency of senility towards close mindedness and false certainty. Age related cognitive decline tends self-imposed filter bubbles, because our own need for certainty overrules the irritating doubt that comes with continually throwing oneself against the endlessly variegated multiplicity of the world. The social extension of a filter bubble, so called online “echo chambers”, are warped epistemic communities in which confirmation bias runs rampant and each member of the community is selected and belongs to it by virtue of the fact that they bring information which aligns with the prevailing consensus.

While the argumentative theory of reasoning helps explain how some cognitive biases is strategic ignorance it is less clear if all biases fall under this header. Other types of bias appear to be distortions of memory, statistical reasoning, or pattern recognition that indeed fit the traditional view of erroneousness plain and simple without any implicit, self-centered benefit. This faultiness is explainable if the brain is not a fine tuned general purpose problem solver but an improvised collection of jerry-rigged, unoptimized subsystems. Borrowing a term from software engineering, Gary Marcus describes the brain as a kluge, a slapdash, patchwork assemblage of half-measures and messy circuitries that loosely hang together but nevertheless manage to get the job done (Marcus, 2008). The reason for this, Marcus argues, is that evolution has no foresight so genes that confer immediate advantages are preferentially selected over genes that promote long term synergies in the genetic makeup. The genes that specify cognitive traits do build up toward a more perfect reasoner, instead over the course of evolution they could even get in the way of each other and accumulate sloppy "legacy code" that consigned the cognitive system to sub-optimum overall function. Seen as a kluge, the brain embodies an engine of ignorance in the full sense of the word as much as it embodies an engine of knowledge.

The salience bias, for example, is when more highly emotionally valanced—such as upsetting—information is prioritized rather than more neutral information even if the latter is more epistemically relevant. Sensationalist news, outrageous claims, and prophecies of gloom and doom stick out further than more muted, if more truthful, statements. Comparing it to an idealized "fleshless" model of a rational agent, salience bias is indeed inexplicable. The emotional significance of a fact should not indicate anything about its truth status. Yet salience bias, like its related negativity bias (the tendency to preferentially remember bad episodes) makes sense from an evolutionary psychology perspective. The emotional centers of the brain are primal decision systems timestamped to a bygone prehuman age that makes split second gut reactions (Nesse and Ellsworth, 2009). Remembering negative events more easily would make sense in a world where deadly consequences have more immediate

(and fatal) significance than happy outcomes. Ideal rationality chafes these underlying brute impulses, but that doesn't mean it is any less rational when viewed in its proper context. In a manner of speaking there is a prehuman subcortical rationality operating underneath and through neocortical human rationality. This sub-rationality follows a primordial logic that lags behind contemporary situations and circumstances and so it seems illogical from the standpoint of the validity criteria that apply to conventional reasonableness. Only later in evolutionary history did the prefrontal cortical structures show up which we now use to reflect on the primal subsystems that stubbornly pull our strings in absurd directions despite our best efforts to see through them.

The organization of neuroanatomy into layers demarcates a series of differentiated echelons of priority of control. Older structures are more strongly implicated in survival and are responsible for the most ancient and embedded of priorities. In general, the closer the proximity to the brain stem the more evolutionary older the organ, and the more central its role in regulating involuntary vital functions. The evolutionarily oldest hindbrain structures that are shared by all vertebrates control autonomic functions such as breathing and heart rate. Wrapping around these structures is the midbrain which controls emotions, memory among other functions, and then finally the forebrain and cerebrum in which all voluntary actions, language, decision making, and other higher functions are mediated.

Generally, the lower the structure is, the more priority it takes over the higher structures, and vice versa, the less control the higher structures have over the lower. Hindbrain structures are "read-only"; it would be a poor design indeed if one had to consciously control one's heartbeat. Similarly, autonomic disturbances such as panic attacks override and dominate the conscious process because of runaway feedback loops of hyper-activation in lower structures such as the amygdala (Sobanski and Wagner, 2017). Such layers of access control and prioritization arguably explain many cognitive biases. The evolutionarily older lower structures have privileged say, thus giving emotional valences overemphasized significance in reasoning processes. According to the society of mind hypothesis, the brain's architecture is composed of relatively "dumb"

components that work together to achieve more robust intelligence (Minsky, 1986 Society of Mind). Normally, these modular systems operate in relative harmony, like a federation of semi-autonomous states or parliament. Sometimes, however, one subsystem can take over unilaterally, causing irrational behavior that emphasizes spurious coincidences, false associations and all the other stuff of contentful, positive ignorance.

The predictive processing model also sheds considerable light on cognitive biases. Sensory states are actively anticipated by feedforward guesses that meet occurrent inputs in the middle. Such a mechanism is susceptible to errors of over and under-prediction. In extremis these errors can completely unmoor perception from reality to the point of pareidolia—the perceptual illusions whereby the subject senses coherent patterns in irrelevant stimuli configurations such as faces in the clouds, religious figures on pieces of toast, or voices in the wind. Such distortions indicate exceeding sensitivity in the feedforward signals, such that the predicted image overwhelms the contextual specificity of they incoming signal.

In broad terms, several categories of cognitive bias can be identified: social, decisional/belief, memorative, and pattern-recognitional. Generalizing over this variety is no small task, but the idea that there is either some implicit self-serving benefit in bias, or it is explainable in terms of the kluge-like slapdash construction of the brain and prediction errors, goes a long way toward achieving it. Cognitive bias is a kind of low-level, preconscious, strategic but not quite willful, ignorance. It is structural ignorance in the most comprehensive sense in that it structurally determines that ignorant conclusions are generated by the cognitive engine.

This concludes the part of this section concerned with characterizing the brain as the "organ of ignorance". To get a fuller picture of naturalized agnoiolgy we must go beyond the brain. We must even go beyond systems that are not conventionally thought of as epistemic concerns. While the mind and brain are natural units of epistemological analysis, there is a deeper informational level at which ignorance—broadly construed— and uncertainty display their shaping effects on complex systems. This more comprehensive approach allows the true

context of cognitive ignorance to be properly situated.

If adaptation to the unexpected is clearly identifiable in the connective dynamics of the brain, it must also be evidenced by biological evolution. In each case an openness and flexibility is counterposed to a setting of environmental volatility is what shapes and molds the adaptive system. The evolution of nervous systems is itself a solution contained within this broader compass of adaptiveness implemented by natural selection.

Several mechanisms for evolution are widely accepted by the scientific community. Among them are random mutation, genetic drift, and differential reproduction. Together these define descent with modification which leads to speciation and biodiversity. Evolution as a whole can be summarized as change within a lineage in response to external changes in the environment that preserve the lineage's ability to generationally replicate a version of itself modified to incorporate those changes (adaptations). Evolution is therefore an information system responding to spikes in the unexpected much like the the model of an epistemic agent described earlier. There are generally two major views of evolutionary process, the more conventionally accepted gradualism versus catastrophism. The original view of darwinian evolution is that decent with modification was a gradual, more or less constant progression. The sudden, massive efflorescences in biodiversity seen in certain geological epochs, such as the Cambrian explosion, were attributed to the incompleteness of the geological record rather than to some causal factor at play. Eldredge and Gould later proposed an influential theory called punctuated equilibrium that contested this gradualist perspective (Eldredge, Niles and S. J. Gould (1972). "Punctuated equilibria: an alternative to phyletic gradualism" In T.J.M. Schopf, ed., Models in Paleobiology.) Punctuated equilibrium states that evolution cycles between long periods of plateaued stasis and punctuated catastrophe, at which point heavy selection pressures are placed on populations and there is a spike in the branching of species called cladogenesis. The sudden spikes of evolutionary change in the fossil record are therefore reflective of an underlying cause rather than spottiness. The explanation for this is that period of environmental stability

do not pressure require adaptations, as adaptation is a response to a changing environment. Extremely disruptive events drive speciation. In a way that is exactly analogous to the work cycle that describes learning and uncertainty reduction in cognitive systems, certainty results in stagnation, a stable environment that inhibits, or rather fails to elicit, activity. Equilibrium begets equilibrium, the present is the same as the past, the new the same as the old. It is surprises which spur change, sudden "wake up calls" that spring the mechanisms of evolutionary innovation into action.

Evolution is also marked by random mutations, most of which will disadvantage the organisms that inherit them but some of which will impart beneficial traits that lead to the differential reproduction of some lineal branches over others. The notion recalls the idea that there is richness in sufficient randomness. In line with the law of large numbers, random, or near random DNA replication errors over the course of countless trillions of instances, will eventually produce a "beneficial error." The system's own imperfection, and its inability to detect and correct it, ultimately leads to its continuation as an uninterrupted chemical reaction over the course of millennia. Indeed, the flawed replication of DNA is one thing that distinguishes it from inorganic matter like a crystal. Crystallization is a process of perfectly periodic cumulative repetition. Anticipating the discovery of DNA, in his *What is Life?* physicist Erwin Schrödinger posited the existence of "aperiodic crystals" as the carrier of hereditary information, to underscore the idea that the essential aspect of the genome is that it is able to encode irregularities in its structure and therefore express varied unique combinations of traits.

In the terminology of the Peircian theory of signs, a biological adaptation is an *index* of an environmental change or dependency. Like a weathervane pointing in the direction the wind is blowing or a bullet casing in a crime scene indicating a gun was fired, an adaptation is a reference to an environmental feature. For this reason the shape an animal's teeth, either flat molars or sharp canines, predict either a herbivorous or carnivorous diet, or in the case of a species with mixed dentition, omnivorousness. An index connects two parts of the world, registering information about its referent indirectly through the sign. Because each adaptive

feature of the organism is such an index, the whole organism can be conceived as a kind of formal system: the semantic correspondence between an adaptation and a component of a ecological niche functions as a symbolism by decomposing into a set of semantic ordered pairs. Each adaptation is a vector pointing to a niche. As Dennett puts it, "Any functioning structure carries implicit information about the information about the environment in which its function 'works.'" (Dennett, 1995).

This last point about the symbolization of biological traits is important because it counters probably the main criticism against viewing natural selection in epistemic terms: namely that it is not an epistemic agent because it lacks semantics. How can knowledge or the lack thereof apply to what is an otherwise mindless process? It would be like attributing epistemic states to the weather, and wondering if a cloudy day will “make up its mind” and “decide to rain today”—a glaring category mistake. This is not the case however if evolution is viewed as an information system, a view which, as argued previously, applies to the whole physical universe. Information is the raw ingredient of knowledge. You can have information without knowledge but not knowledge without information. The one thing that is missing, and which keeps evolution from being explicitly interpretable as a knowledge process, is its lack of a subjecthood to self-reflect on its own problem-solving behavior and therefore dub its activities as knowledge-seeking. Thankfully, we are more than happy to step in and take on that role and to reflect on evolution and see that it is indeed a problem solving mechanism. When an organism is seen as a symbolism, as a kind of consistent and true statement written in a robust genetic programming language, evolution is directly tied to a semantics and this designates it as an epistemic concern. Indeed the fact that such a language can be read into biology is what makes taxonomic classification so eminently possible in biosystematics. Genes and the traits they express form repeating, stable “alphabetic” or “dictionary” units that reappear in an organized pattern which can be reconstructed as a phylogenetic “narrative”. It is no coincidence that both grammar and phylogenetics are both describable as treelike data structures. There’s such a semantic clarity across different taxa that an underlying syntactic unity can be perceived between

animals as far flung as horses and tiny, inch long seafaring lancelets. Given both taxa share the "morphemes" which define chordata : the presence of notochords, dorsal nerve cords, pharyngeal slits, and tails. (And much like a high level programming language, these recombining phenotypic traits can be further compiled into a genetic "machine language" of 1s and 0s, or as the case may be, nucleotides.)

The matter is further clarified when we consider the teleonomic properties of biological systems. They have, or are at least perfectly describable as, goal directed, intentional agents. To the extent that evolution is an algorithm, it is as much an epistemic concern as any belief or mental state that implement algorithms. Evolution is computable.

As Dennett describes it, evolution is a kind of search algorithm that iterates over a "design space" of all biological possibilities. Biological possibility is in turn constrained by physical and logical possibility and is extended by historical possibility. For historical possibility there are certain adjacent state between items in design space which are more accessible given the current state. They fall within the scope of the same local minima. Therefore the expansion and brachiation of the tree of life from a given earliest ancestor, describes a series of path dependencies that further constrain evolution down given predetermined bottlenecks. Natural selection provides a mechanism for pruning and preserving elements of this tree such that the entire process is able to recur indefinitely, keyed to environmental changes as they arise. (It's worthwhile to note that the "environment" can contain other species, hence coevolution, and point that evolution can be endogenous, that is, entirely self-causing as one part of the tree pressures another part as with predator-prey antagonisms.) One may notice a sharp similarity between design space and the concept of discovery space detailed in chapter 8. This is no superficial resemblance. It requires a simple perspective shift to redescribe design space as discovery space. Should we chose to adopt the intentional stance, evolution can be spoken of as searching through a landscape of ignorance and arriving at knowledge states that reveal whole parts of the landscape from which the agent can continue to proceed. While the exploration of discovery space is agentive and purposive, evolution, viewed as

an algorithmic process, is perfectly definable in these terms. On this analogy, implicit knowledge is embodied in the design structure that builds up over geological time.

There is another sense in which evolution can be seen as a product of ignorance. It is blind. It produces knowledge as design structure while itself being completely ignorant. Again such an interpretation might be viewed as a category mistake, bun insofar as an entity has a capacity to learn it has epistemic status (DeNicola, 2017). Evolution is therefore ideally suited as a negatively defined, eliminative learning process that yields residues of viable design. Evolution arrives at the viable by boiling away the inviable.

Evolution is also relevant to the study of ignorance for the obvious reason that it is the process that shapes the brain. Evolutionary psychology explores how cognitive systems are evolved with various ins and outs that are keyed to evolutionary history. Such studies also unveil how certain evolutionary design pathways set us up for various forms of folly, bias, and cognitive distortions. Human beings are not optimized to be perfect reasoners but to think in ways that benefit their biological fitness.

One more example will suffice to drive home the argument that ignorance as objective uncertainty manifests on multiple levels. While secondary to neural growth processes and evolution by natural selection, the immune system is an important biological network that learns by accumulating effective defensive strategies over time. As a system capable of undergoing learning, it too shows signs of uncertainty reduction characteristic of epistemic agencies.

Any system that can be said to learn shares a set of defining traits. It must respond to environmental novelties by formulating a new persistent memory structure that encodes information about the event that caused it. A learning system must also have a discriminator: a means of recognizing that a specific environmental change has surprisal value and is a "difference that makes a difference" which is relevant to its process. Learning is therefore the characteristic response to ignorance. Wherever learning happens ignorance precedes it. A system that learns therefore has epistemic status. Logically, because

no finite thing can be omniscient, to be a knower means to be able to learn, and to learn presumes ignorance. Ignorance is the sea in which knowers swim. All three parameters, knowledge, learning, and ignorance are required to specify a knower. You cannot specify a knower without reference to ignorance in the same way you cannot specify the flight of an airplane without reference to the medium of air. To be a knower means to have been learner, which means to have at one point been ignorant. Any state of knowledge is a far from equilibrium state that was preceded by uniform ignorance in every case and which is surrounded and contained by it. Learning is equivalent to the energy inputted in order to push the state of knowledge out of equilibrium into a more ordered state.

Such mnemonic systems can be found throughout the living world, most notably as we have seen in the case of neural networks and evolution itself. These sub-agentive systems may not immediately present as things that learn because we tend to associate learning with rational agency. Yet on close inspection they fit the definition flawlessly. A genome and the set of adaptive functions it maps to comprise a collection of semantic indexes that represent a built up memory of environmental changes that is genetically persisted (albeit with some mutative volatility) through the generations. Importantly, a learning system does not need to have subjective states to be considered an epistemic agent. Having subjective states and being able to self-reflect on one's ignorance is doubtlessly a more advanced feature of a learning system, one that plays a role in proactive (knowledge-seeking) versus merely reactive (passive, reinforcement based) learning. Yet the strict requirement of subjective states would lead to absurd rejections of the part which forms the whole. The neural connections that form as part of Hebbian associationism and which physiologically underserve the experience of learning, would otherwise not be counted as true learning because the parts that undergo it do not by themselves have the subjective experience of the whole. One might qualify it to say that subjective states endow knowledge with ultimacy; such knowledge states are totally realized while non-conscious knowledge states are rudimentary or partial. The latter are not quite knowledge states, but rather information states. A whole human being is indeed a more robust epistemic agent than a mere "autonomic" learning system. Yet

knowledge can be tacit and implicit: such knowledge is implicitly embodied in the information structures that accumulate in a physical system of sufficient complexity and organization. It is all but knowledge in the ultimate sense of being bound to a self-aware knower. If knowledge, in other words, can be said to be instantiated in a particular configuration of brain mater, then it can also be instantiated in an otherwise insensate network of information processing and memory encoding components. Such functionality does not need to be an intentional object in the mind of a knower to still qualify as valid acumen. In any case, if we are to maintain the existence of tacit knowledge these information states must also be respected as epistemic concerns. I don't know how I know how I am conscious, but it is clear on some level that my brain "knows" how, as proven by its capacity to produce it.

For the sake of a more rounded argument a third naturalistic learning architecture ought to be described. Immune systems occupy a compelling middle ground between evolution and cognition. The goal here is not to give a detailed scientific description of the immune system—a task which goes well beyond the horizon of my own discovery space—but to analyze it particularly within the inverted epistemic framework. As an information agent, the immune system is tasked with collapsing uncertainty, no more and no less than the mind or brain of a traditional knower. Because thinking about such natural systems in these terms is so unfamiliar it promises to offer some unusual insights.

The immune system exhibits a suite of processes of dizzying complexity that effectively amounts to a kind of "proto-intelligence". Consisting of a genetically pre-specified, fast-acting innate primary immune response and an acquired, slow-acting adaptive secondary response, the immune system is the body's armed forces in the endless conflict of the tree of life's civil war. Its two-tiered, fast and slow organization is evolution's way of covering all its bases at every timescale. Various ancient enemies, persistent microbial and viral threats for many generations, exerted such pressure on the survivability of a population that only those with a genetic composition that provided natural resistance to the invaders survived. This formed the foundation of the primary immune system. The immune system warehouses a stockpile of domain general antibodies that

can respond to a broad class of pathogens. Because of the vast populations and rapid reproduction rates of pathogenic microorganisms, however, immunological countermeasures must also provide mechanisms to mount targeted, custom defenses. Hence the secondary immune system's ability to construct and retain ad hoc antibodies in the event of an evolutionarily novel infection. Thus there is a generic and fast and a specific and slow immune response. The acquired immune system even displays proactive, tacit knowledge of its own incomplete information state (ignorance) by secreting antibodies into the blood stream that constantly "patrol" and keep watch for pathogens (Burton, D. Antibodies, viruses and vaccines. Nat Rev Immunol 2, 706–713 (2002). https://doi.org/10.1038/nri891). Such exploratory behavior is what one might expect from an epistemic agency that seeks to map its own ignorance. The immune system implicitly behaves as though it must explore discovery space to keep one step ahead of pathogens attempting to crack and breach its security protocols.

The immune system contains features that fit the description of a cognitive agent. It has sensors, called pattern recognition receptors, that detect specific molecular features of pathogens or cell damage indicative of infection (Kumar H, Kawai T, Akira S (February 2011). "Pathogen recognition by the innate immune system". International Reviews of Immunology. 30 (1): 16–34. doi:10.3109/08830185.2010.529976. PMID 21235323. S2CID 42000671.) Thompson MR, Kaminski JJ, Kurt-Jones EA, Fitzgerald KA (June 2011). "Pattern recognition receptors and the innate immune response to viral infection". Viruses. 3 (6): 920–40. doi:10.3390/v3060920. PMC 3186011. PMID 21994762.) It has mechanisms of "attention" such as inflammation, which causes certain biomarkers to be concentrated in an infected area that attract white blood cells by increasing blood flow. And it has a memory. Some cells that were generated in response to a specific pathogen are retained over time, thus expanding the immune system's repertoire over the long term. Everything points to the immune system having a sort of collective, emergent, swarm intelligence. This is the intelligence of the ant hill. It's worth stressing that this is probably the same sort of principle of intelligence working within us: each neuron is an ant in the brain hive.

As with brains and evolution writ large the immune system, with all its related cellular signalling networks, nodules, antibodies and biomarkers, is an epistemic agent that gains in structure by repeated confrontation with its own ignorance. Ignorance here means information incompleteness. No completed brain or immune system, or intelligence system generally can be initiated from the top down "fully stocked" with a total repetoire of structures fit for every possible circumstance it might encounter in the course of operation. The objects —pathogens in this case— making their debut in those circumstances may not have even existed at the time the system is initialized. Any finite physical system, taken as closed and separate from the rest of the universe, is necessarily poised at a deficient information asymmetry from the larger portion. More can happen than has happened; this is the ontological principle that assumes adaptation and intelligence in the first place. A learning system works to close that gap and reach a limit that it can never quite reach but seeks anyway: equilibrium, information symmetry between it and its environment. To do so it must acquire content related to facts about its environment that are impossible for it to possess initially. Thus knowledge is "driven through" ignorance by learning, as uncertainty collapses and its knowledge base and memory crystalizes.

What do the entwined threads of biological evolution, immunology, and brain networks tell us about the naturalistic origins of intelligence? It shows that in a perfectly simple, certain world, there would be no intelligence. Learning and intelligence are contingency plans for the complex and therefore uncertain. They represent a flexible openness which stands in relation to, and tacitly "expects the unexpected", so that the biological system can meet the unknown at multiple spatiotemporal scales. Biological complexity is an adaptation to environmental complexity. Selectionist and adaptational principles in brain development, immunology, and evolution embody trial-error feedback, error correction, and control processes which define a primordial relationship with ignorance.

Here the relationship between intelligence and complexity becomes clear. Physical systems with highly repetitious, stable patterns such as crystals have low entropy. Systems with more structural diversity have higher entropy. The

Kolmogorov complexity of an object defines the minimum length of a program required to yield the object as output. A repetitious objects can be described with a short program while a complex, highly differentiated object needs a longer one. In a closely analogous way, to think is to formulate an algorithmic description of the complexity of an object. It is to simplify a problem by compressing the information about into the most generic, reusable form. We have to think harder about complex problems. Intelligence therefore is an adaptation to complexity just as much as it is a property of complexity. A simple unicellular organism that lives in a puddle deals with a simple environment, consisting exclusively of creatures and debris of comparable or greater simplicity. Thus its behavioral rules are simple. Its ecological niche locks it into a cycle of interactions with its world that demand no innovations and which genetically prescribed reflex actions will suffice for it to navigate. For higher animals their adaptations are a response to a more complex environment — their "umwelt" is richer and deeper, fraught with more varied threats and replete with more varied opportunties. Their intelligence generally speaking scales with this rise in complexity.

Another way of putting this is that higher animals' environments tend to be more complex *because* they are more intelligent than a bacterium. Enhanced perception and sensorimotor abilities "brings out" the latent complexity of the environment which for a simpler organism would simply be imperceptible. To the more ignorant, or rather to the less capable of knowledge, everything is simpler that it truly is, "dumbed down" as we say. Both ways of putting it are equally true and comprise both sides of the same evolutionary feedback loop.

In a sense simpler, "intelligenceless" or sub-intelligent as opposed to "dumb" organisms cannot be said to be ignorant or unintelligent; as they have no need for knowledge and thus have no deficit in it. All their capabilities are embodied by their genes and in their somatic rote operations. Such organisms lack the potential for knowledge and are thus their behavior is purely a function of their morphology. For the human brain, the matter is much different. The brain is in all likelihood a universal computer. Our umwelt scales to the entire universe of possibility and as such we are confronted with maximum complexity. There is much more we are capable of noticing. This is in no small part explained by the

fact that during human phylogenesis our own kind, with equally complex brains and behavioral repertoires, comprised an at times evolutionary cooperative or adversarial game spurring us to constantly synergize or outwit each other and thus driving up the computational challenge level along the way. The human brain's operational environment is therefore as maximally uncertain as possible. All change is "blown open" to it. The ultimate purpose of the mind is then to contract this entropy, to locate and representationally construct stable pockets of order. It is to identify chains of regular subsequences within the overall randomness of the pre-interpreted universe.

Intelligence increases as the perception of entropy increases. In a perfectly ordered world, like in a perfectly simple world, intelligence would not be necessary. Certainty begets triviality. The more perceptible entropic complication in a system, the more processing power is necessary to surmount it. Thinking can be conceived as "possibility mining", that is, as the search for the actual within the possible or merely probable. A dynamic, complex phenomenon implies an overflow of possible configuration states, which must be sorted through by the ordering powers of cognition. Possibility is weighed against sedimented preconceptions, so that what fits with the prior is admitted as the most likely scenario. With possibility comes uncertainty and with uncertainty thought.

Complex functional systems deal in information. Information is that which eliminates excess irrelevant parameters in the course of maximizing a utility function, leaving only kernel of representational value. Conversely, uncertainty and ignorance scale with complexity. Understanding the relationship between brains, evolution, and immune systems—and complex biological systems in general—helps underscore how the complexity of the world fixes uncertainty into place as a universal condition to which all life must adapt.

It's only too fitting to close this chapter on naturalized ignorance to apply the principles of constructive ignorance to nature itself. In a provocative 2021 (https://arxiv.org/pdf/2104.03902.pdf) position paper, an international team of researchers proposed what they call the Autodidactic Universe. It's a theory of cosmology inspired by unsupervised neural network machine learning models

that describes the universe as selecting its own laws from a set of possible laws. The central mystery that they seek to take on is why the the universe exhibits one set of laws despite that according to modern physics there being a large number of other equally consistent sets? They explore an evolutionary perspective on cosmological laws in which they are not eternal and immutable but changeable. This notion of change however is of a greater order of depth than that of natural selection, which as described previously exhibits the most rudimentary trial and error ignorance-reaction logic. Similar to the definition espoused here, they propose a more nuanced definition of learning which encompasses natural systems more broadly. "We can say then that such a system learns when it is able to alter its internal processes and actions in the world to better capture and exploit the flows of energy through or near it to further a goal, which is typically continued existence, but might become more than that." (p.7) In short, learning is the accumulation of adaptive structure.

The authors come tantalizingly close to an ignorance-first understanding of learning. "The behavior of a learning machine may include steps where the degrees of freedom are thermalized or quenched. (p. 11)" This is another way of saying uncertainty is reduced, the contraction of probable alternatives to one.

The details of this theory are quite involved. What's of relevance is the philosophical and negative-epistemic implications of this vision of the universe as an autodidactic learning system. From one angle this is a *vindication* of the *objectivity of ignorance*, given that as I maintain all things to which learning may be ascribed are subject to the principles of ignorance. The universe is everything that physically exists, therefore to be counted as something that learns it must respond to some negative principle which nevertheless ontologically informs objective conditions.

One issue is whether learning requires agency. Agency requires a certain functional unity and purpose. "The universe" might be categorizable as a unity, but it is equally conceivable as simply a collection of all that physically exists and their relations. One might add the regularities (laws) that bind its objects, give them properties, and consistently predict their behavior. Yet one could equally get by saying that it is the dynamism between the elements of the collection,

rather than the whole of it as a unit, which define the laws. Is universal law deductive, going from the whole to the part, the class to the instance, or inductive—the opposite? That is, would it be a reification to assume that the laws are themselves existent information structures somehow encoded into reality, the DNA of the cosmos, rather than simply how things just happen to settle and sort themselves out without any other substrate to scaffold that incidental arrangement? The autodidactic universe paper assumes that possible physical laws exist in the nominal sense that the universe explores them in the course of its time evolution. If so then where are they? The only answer is that they exist in ignorance as a negative objective substance.

If the universe is a learning system and operates as a coherent unit then it is a priori true that it is an unsupervised learner, as there is nothing outside of it to teach it.

The anthropomorphic principle which is often trotted out to explain away why the universe is so seemingly fine-tuned such as to support intelligent life because there is such life to observe the fact, is so unsatisfying because it commits to a basic circularity. The autodidactic universe theory seeks to provide a more robust explanation for why the universe's parameters are so tightly adjusted to each other, but there are many problems with this.

The greatest flaw with the argument is that it does not account for the role of entropy in information dynamics. The universe is an energetically closed system —by all accounts *the* closed system—no more energy flows into or out of it than was already present as per the law of conservation. According to the theory supported here, however, learning is precisely associated with differential energy flows, as a gradient in energy also specifies a gradient in information. Therefore only open systems are capable of undergoing learning.

An additional problem with this theory is that there must be laws prior to the evolution of natural law to inform the context of that evolution. If the universe is a learning system there must at first be implicit rules governing its learning rules. If that is the case, then the true, unevolved natural laws are already present prior to the fact. These are the laws which describe the universe's learning mechanisms rather than its final results. Otherwise, an infinite regress enters the picture in

which there are no originative initial conditions which bound all physical possibilities. The learning rules would themselves have to be learned, and so on to infinity. There would have to be meta-laws that govern the evolution of physical law, but in that case, it might as well be that the meta-laws are the actual natural laws.

I argue that complexity is a necessary concomitant for learning. Since the universe contains all that is complex, it is the most complex thing of all and necessarily more complex than any of its elements. Once more, it is through a habit of projective attribution or linguistic classification that we consider the universe as more than the sum of its parts, to be thing itself rather than just the collection of all things. In which case there is no "the universe" only a set of processes and events which all share the class property of "this-worldness". In which case the universe can contain systems or be composed of systems without being one itself. It is merely a stage, or more precisely, it is nothing more than the objects, properties and relations between all that physically exists. The paper elaborates on its concept of systems:

> To proceed to define learning we need to specify some terms: We speak of "subsystems" and "processes". A system is an identifiable part of the universe that contains a number of processes and persists over a time scale that is long compared to the timescales of its constituent processes - and is often characterized by a boundary which isolates it from processes not contained in it. What is not in the system is called the system's environment.
>
> A system has a set of activities carried out in time, by means of which it does work: it catalyzes or alters itself and other things in the environment. Each of these is called a process, and each has a typical time scale. Each process is closed, which means that it operates in cycles that bring it back to its initial state, after having done work. The thing that a process does that furthers the whole of which it is a part, is called its function. We require that a system must be decomposable into a network of such processes which are interdependent on each other, in the sense that the continuation, and growth of the whole is dependent on and furthered by that of the processes. This is sometimes called a "Kantian whole".

> Typical systems of this kind are found far from thermal equilibrium, where some of their processes channel and control the rates of energy and materials through them. They develop from the need of the whole to manage flows of energy and materials through it. We can say then that such a system learns when it is able to alter its internal processes and actions in the world to better capture and exploit the flows of energy through or near it to further a goal, which is typically continued existence, but might become more than that. (p. 7)

These criteria however are not in my opinion sufficient to describe the universe as a system. The universe does not "happen in time" it is time. It has no functional linkages between the processes that compose it, they just *are*. The universe just is. It goes on to say that there are a class of autodidactic learning systems which "construct its own criteria." Another term for this is autopoiesis: a system which maintains and self-perpetuates its own structural integrity and whose teleological processes are all specialized towards repeating its structure. Biological systems are the most obvious example of autopoietic type, with survival and reproduction being the regulatory background criteria by which organisms adjust their functions and by which they persist their state. The universe however is not trying to maximize any utility function, and this is another flaw of the autodidactic universe thesis. It assumes that the laws of nature are not arbitrary, but there would therefore need to be a reason why one set of laws was selected over another.

What then is the minimal possible learning system? Previously it was argued that natural selection is a learning process, even though by the same definitions natural selection *just happens*. However, biological evolution results in the accretion of functional organization in more directly palpable learning systems—lifeforms. In my view, learning only happens within parts of the universe in relation to others. The paper goes so far as to say that *gravity* is a learning system. Yet gravity as it is currently understood is purely *relational*. Like the misleading noun "the universe" suggesting a thing that isn't a coherent unit, so too is it misleading to think of gravity as anything but the masses which generate it through their interaction with spacetime.

The theory put forth in this monograph about cosmic autodidacticism is

extremely nuanced and sophisticated. I make no claims of canvassing its argument completely. However for something to be self-learning there must be a self, a bounded system of mutually interconnected parts, to which it can refer. The argument strives to describe the universe as analogous to a machine learning algorithm, but for that to be true it must be a machine. The main purpose for bringing up the theory is to examine how it relates to the notion of objective ignorance.

Part V : Homo insapiens: The All-Too-Human Implications and Applications of the Theory of Ignorance

We have now reached the point in the exposition where the theory of ignorance has more or less been articulated. What's still lacking is a more fleshed out series of examples and demonstrations of these principles. All the better to animate and give life to the abstract, metallic theoretical considerations which have filled the discussion up till now. There are many pertinent, practical, social, and psychological real world cases to which the above principles can be applied. Ignorance is alive and well in the world today, all the more reason to marshal and deploy the tools at our disposal to see through it and bring it to heel. Because the ideas in part V are more illustrative than explanatory, I have decided to separate them from the main thread of the exposition and place them in a separate catch-all grouping. This part of the book is therefore more eclectic and less organized than the previous. Each subsection is generally self-contained with respect to the others, and they do not necessarily follow in a linear narrative sequence. The goal is to give a "dragonfly-eyed" wide-angled snapshot of the many situations, conundrums, and real world experiences in which ignorance is a driving force while making use of what we have learned.

Ignorance is a fact of nature. Just as we are caused to know on a most basic level by the fact that the world has a definite reality which impresses itself upon us and so makes certain actions either congruent or incongruent with it, and just as reasoning about this objectivity eliminates informational degrees of freedom

until uncertainty falls within expected bounds, so too are we caused not to know by basic facts. Much of this study has occupied itself with the logical essence of this substantial ignorance. Or it has concerned itself with a low level of description of the information processing of said uncertainty. What has been given less attention is the purely human side of ignorance. To err is human, and being no strangers to ignorance, understanding people's relationship with it who after all seem to have a special affinity for it is indispensable.

The human species name, *Homo sapiens*, "the wise man" somewhat self-flatteringly puts our intellectual aptitudes at the forefront. Yet as the capacity for knowledge increases so does ignorance. The more you could know the more you don't know. A rock is not less knowledgeable or more ignorant than a person. Its epistemic status is undefined. It is unable to experience either state. So in not too cheeky of a fashion, humanity can also be dubbed *Homo insapiens*, because if anything what humankind does not know has been just as constitutive of its historical odyssey as what it does know.

Ignorance Generation

There are kinds of things that put more ignorance into the world. Just as there are factories and smokestacks that pump more pollution into the atmosphere, so there are practices, processes, technologies, customs, and ways of thinking that generate ignorance. There ought to be a word for such ignorance generation, so I propose the term *agniogenesis*. DeNicola identifies 5 types of what he calls constructed ignorance, or deliberate nescience. These five categories are rational ignorance, the decision not to know something judged to be not worth learning. Strategic ignorance, is used to gain a social advantage over another party such as deniability in legal proceedings. Willful ignorance, which in the simplest terms is the intentional act of blocking out truths one cannot handle or which would muddy one's thoughts. Privacy and secrecy, which concerns the regulation of a knowledge boundary between individuals. And lastly forbidden knowledge, knowledge that is socially proscribed against and labelled taboo.

The question is: precisely what mechanisms lead to the construction of

ignorance? Given that ignorance is typically viewed as a lack of knowledge, and given that this is its most common, perceptible manifestation, it is not clear how a lack of something could be created, manufactured, or otherwise invented. It would be like making emptiness more empty As will be shown however, defining the construction of ignorance simply as a negative concept, as the withholding or destruction of knowledge, is insufficient.

If ignorance is misunderstood as the mere privation of knowledge it doesn't make sense that something can make more of nothing. It's like adding a negative to a positive integer: you should only get less of a something. An amnesia inducing device like the neuralyzer from the science fiction movie *Men in Black*, which flashes a light that causes forgetfulness in a subject by wiping their memory, might count as such an knowledge subtractor in its function as a memory eraser. Forgetting is the removal of a memory and is a way that knowledge passes into ignorance via the loss of neurally encoded information. In this, albeit fictional case, there is literally a physical device that causes ignorance. Similarly when the Library of Alexandria burned to the ground ignorance could be said to be produced because knowledge, or rather literary pathways to knowledge- were destroyed. Just as everywhere energy goes from a useable to an unusable state through thermal dissipation, so does information decay and degrade over time unless otherwise backed up, duplicated or perserved. It is no contradiction to speak of a negative as increasing, even beyond the scope of any positive reference value. Certainly this makes sense in mathematics. As DeNicola (2018) observes, "It is true that we bestow an intuitive primacy on positive concepts and affirmations; their correlative negative concepts and denials seem wholly parasitic on them. But not all conceptual polarities are fixed; many may be reversed, as with a figure-ground reversal. We can view silence as the absence of sound, or sound as the absence of silence, depending on context and on our desires and purposes. The same is true of ignorance and knowledge." Given these considerations it is not at all unsound to think of ignorance as something capable of undergoing increase and production even without a corresponding decrease and negation of knowledge. Especially considering that ignorance can have content, as with a misguided attitude, a false belief, or mistake one does not

realize until after it is committed. Such things are not mere negations of true beliefs or mental content. Ignorance as it were "got there first", it does not debase preexisting true knowledge.

So ignorance can increase without a corresponding decrease in knowledge. This increase is not necessarily subtractive. When a false rumor spreads it does not remove knowledge of the truth, that knowledge was never present. Most likely there was no opinion on the matter at all. The rumor is primary and set the tone for what people believe. The rumor specifies a new content which one was not previously aware of and sparks controversy over a misbelief. The phrase "rumor mill", a group of people that spread unsubstantiated hearsay that "muddies the waters" is telling. It posits the image of a social infrastructure for the generation of ignorance. Specific cultural, communicational, and cognitive activities manufacture positive ignorance. Whether it be lies, biased reporting or "fake news"; whether it be misinformation, photos taken at misleading angles, unverified rumors, fabricated documents, forged signatures, espionage, high tech "deep fakes", camouflage, coverups, superstitions, encryptions and ciphers that obscure data, decoys, fraud, obscurantist speech or writing, pseudoscience, psychological rationalization or repression, online phishing scams, impersonations or just plain stupidity, there's a plethora of things which either cause or add to the general sum of ignorance for many deliberate reasons. Or for no reason at all besides folly or accident.

Importantly these angiogenic processes and artifacts are not merely caused by ignorance, as when error begets error, although that is a live possibility. They can often be the consequence of knowledge, as with a deliberate political disinformation campaign, a clever magic trick in which the performer hides the truth about it from the audience, parents spinning a tall tale to their children about tooth fairies or Santa Claus, or a cynical ripoff perpetrated by a sly and cunning con artist. Such regulations of knowledge boundaries are everywhere. These manipulations often make use of knowledge gradients, imbalances in knowledge and ignorance between two or more parties. Just as knowledge flows from those at a higher gradient (more knowledgeable) to a lower gradient (more ignorant) as from teachers to students, so too can those "in the know"

manipulate the less informed by controlling both sides of the boundary. For the same reason, as cognitive studies of deception have shown, the old saying "you can't kid a kidder" holds true. Someone who is good at lying is skilled at detecting another liar (Wright, Berry, & Bird, 2012). Equilibrium in knowledge gradients, in other words, makes deception infeasible. One has to be in the know and and the other out of it for the deception to succeed.

It's intuitive to view the quantity of ignorance in the world as fixed. The quantity of knowledge varies, but the amount of ignorance is absolute even if it can't be enumerated. It's somewhat analogous to the problem of future contingents. Suppose that I do not know what will happen tomorrow. And suppose I never at any time in the past had this knowledge and forgot it. I did not know what will happen tomorrow three days ago, three years ago, or at any time in the past. My ignorance of the events of this particular date does not vary until I eventually learn what will happen. Ignorance has this inertial quality, where only knowledge changes relative to it but it remains static. Everything that will one day be known was once knowledge existing *in potentia* within the fixed manifold of ignorance, and that ignorance was always "there" since the beginning of time.

This assumption of the fixed totality of ignorance is flawed, however, when we consider that ignorance is tied to events in time and to new developments in the world. One cannot be ignorant of what does not exist, because what does not exist is not a fact. It goes back to James' notion of the unfinished universe, a world that is growing at its leading edge and where new sources of information are produced. The key here is discoverability: ignorance that comes into the view of the horizon and has the potential to go from unknown unknown to known unknown. Someone who was in a coma yesterday and knows nothing about what happened does not have the same quality of ignorance as someone who was ignorant about that same date 1000 years ago, even if they are equivalently oblivious of it on one level. The availability of information is different between them. For the amnesiac, yesterday has become one with today through the flow of time and causality, and thus there is some discernible trace of yesterday in today in the sense that everything that happened yesterday is contiguous with

everything that happens today. The past embeds information traces into the present. The amnesiac can reconstruct yesterday, or at least some part of it. What's important is the *causal role* a particular blob of ignorance plays in its relational interaction with everything else that's apropos to the matter at hand.

Agency has an important role to play in the generation of ignorance. Ignorance *can be caused*, it is not merely a passive state. Consider the case of someone who is viewing a forged painting who does not know it is a forgery. Are they the same sort of ignorant as someone who does not even know of the painting's existence or who doesn't even know what a painting is, let alone that this one is forged? In the case of the viewer of the forgery, the viewer has been *duped*, they have been hoodwinked into believing that they are viewing a real painting. Meanwhile the one who knows nothing of the painting cannot have this impression created in them. An act has been committed which has made a pocket of ignorance in the viewer but which has not been created in the non-viewer. In this instance, there is an additional ignorance betrayed by the forgery. The difference is a matter of intent. The forger commits an intentional act to deceive the viewer. Something new has happened that is now a going concern, namely the production of a fake posing as the genuine article. One might say that the person who does not know of the painting is simply doubly ignorant, because they would have to discover the painting and then find out it is a forgery. However, they have not had this ignorance inflicted on them as part of a deceptive scheme. It's a subtle but important distinction.

Deception is perhaps the most definitive example of an ignorance causing act, but it's not all. Lying has an intricate logical structure that deals with the interplay and exploitation of the relative knowledge and ignorance between two or more agents and is a rich subject that will be examined in greater depth in the next section. But lying is not the only form of agniogenesis. A mistaken assumption can lead one to behave in ways that create whole webs of ignorance. Once one has committed to a false premise unwittingly, the rest that follows will lead only deeper into ignorance until the compounding contradictions with reality raise suspicions of misbelief. The internet is notorious as a font of conspiracy theories and false claims, which make the world a more ignorant

place. That they are called theories suggest that conspiracies are ordered constructions. Though they may be false, some ingenuity went into connecting those illusory dots. Each claim fits every other claim coherently. In the same way any conclusion can be derived from an argument that is valid (noncontradictory to its premise) but unsound. A dedicated conspiracist can always find ways to interpret anything to be consistent with their theory. Conspiracies are inventive. New ignorance is created, not because knowledge or information is destroyed by them, but because a misrepresentation has been introduced into cultural circulation.

A misrepresentation is a live possibility, * an option to prompt belief*. It is a proposition which renders a picture of affairs that does not obtain in reality, but which causes a mental state that still pictures *something*, a fictional state of affairs that one still responds to and plans against. Believing a misrepresentation rearranges the objects of mental space which would not have gone into that false combination in the absence of that misinformation and prompts actions consistent with its illusion.

Misinformation presents excess options for belief which are like the squid's ink that obscures the truth about the squid's location. It's obscurantist. It clutters and disorients thinking, conjuring red herrings and wild goose chases that throw us off the scent of true knowledge. If information is the reduction of uncertainty, misinformation is the production of uncertainty. It adds more yes or no questions to the list that has to be answered to narrow down the truth yet answering them brings us no closer to it. Worse, it is noise that poses as a signal. Whether to be informed or misinformed is still to be caused to believe. Misinformation, though false and corrupted, still has a form that is homologous with information. Unless it is detected and rejected it is still admitted into the thinking process all the same, in the same way that a malicious software program that bypasses antivirus security checks runs on the computer just as normally as safe code. Once misinformation has infiltrated our thinking and woven its way into its overall fabric, it becomes hard to extract it and to undo all the effects it has ingrained into the entire conceptual framework.

In some cases the tainting of concepts by misinformation is all or nothing,

ruining beliefs altogether like a drop of Botulinum toxin in a glass of water. Once misinformation has wiggled its way into our thinking, so long as it is not quickly caught and expunged, we risk propagating it unintentionally in a chain reaction of uncontrolled ignorance spread. Ironically the widely believed popular myth that people swallow 8 spiders in their sleep every year originated in a 1993 magazine article that stressed how easily people accepted absurd claims online ([https://www.britannica.com/story/do-we-really-swallow-spiders-in-our-sleep].)

As it happens Mark Twain's witticism that "A lie can travel around the world and back again while the truth is lacing up its boots" checks out. Research on the spread of false news on Twitter found that it spreads faster than true news, likely because false news is more novel than true news, and people are more inclined to respond to novelty ([https://science.sciencemag.org/content/359/6380/1146]). This indication that novelty spurred the drive to share false stories is a revealing sign about the intersections of knowledge and ignorance. It shows how, to understand both, they must be understood together, and moreover, process must be emphasized over content. The process that compels knowledge acquisition is the drive for learning, but it would be a mistake to assume that everything that is learned qualifies as true knowledge. On a basic level learning is information uptake, but misinformation and information have isomorphic forms, either can be learned. Miseducation is a genuine risk. Children learn whatever they are taught, true or false. Novelty is what motivates learning. People are less stimulated by information that they already know than by what they don't know. If misinformation is more novel and provocative, then it shows how vulnerable and flawed the knowledge acquisition mechanism is and how the same tool can be applied to both constructive and destructive ends. Conspiracy theories, although they rarely hold up to scrutiny, often present an incendiary and alarming portrayal of events, prompting outrage and morbid curiosity. The mundane truth often cannot stand up to the dramatic power of a rendition of the world filled with angels and demons.

Ignorance production has massive implications for society and for personal epistemic hygiene. Like antibodies in the immune system, the minds defense

against misinformation is critical thinking. In a relaxed, receptive mental state, the gates of the mind are thrown wide open, come what may. This may maximize learning but it also risks the highest likelihood of misinformation infiltrating undetected. Open versus closed mindedness are best as situational policies, states which alternate depending on whether or not certain critical thinking checks have been passed by a given information source. It is mental suicide to simply close the borders of the mind to all outside influences and rigidly commit to one's preconceptions. Mental efficacy is predicated on a constant adaptive coupling with new information and findings in a changing world. Neither can one accept everything one hears with childish naivety. Ideally, the optimal approach is to stop at a new claim—especially if it is provocative and carries a seeming weight of import—evaluate it, and if it passes all rational criteria, then approach it and what it entails with a more open mind.

A major problem with ignorance inducing artifacts is that they are mixed with valid artifacts. In an ideal epistemic utopia, everything one comes across would be valid information and each episode of learning would bring one closer to the truth. One would could breezily embrace all information with an unguarded open mind, feasting on knowledge without precaution. In the gritty, falsehood blemished real world, there is no such easy harmony to be had between the learner and the learnable. Good educational institutions strive for the utopian optimum of a learning environment where everything one learns has been stress tested and pre-approved by disciplined and reputable critical minds. Practices such as peer-review and the citation of sources establish overlapping nets of verification checks to ensure that misinformation gets caught before it can move forward. In the wild frontiers of the internet, one has no such assurances. The much decried rise of "post-truth" politics, climate denialism and other forms of willful ignorance, the tragic ascendency of idiocracy, thrives in the anarchic information wilds of digital culture where it is much harder to validate information and critical standards are nil.

As Hardwig's classic 1985 paper *Epistemic Dependence* argued, knowledge often comes second or third hand. I was never under the microscope to discover that germs cause disease, I owe that one to Louis Pasteur. My knowledge of germ

theory is a matter of trust. It is infeasible to re-derive all the accumulated knowledge that has been culturally produced over the centuries. If everyone who claims to know the results of science had to repeat the experiments for themselves for them to be considered truly knowledgeable, most would know far less than they do now. Trust is thus an irreducible component to efficient knowledge acquisition. Misinformation is therefore particularly insidious because it undermines the trust upon which epistemic dependency rests. The internet is distinctly treacherous in this regard. Idealized as the "information superhighway", the web was originally conceived as a epistemic utopianist vision in which the cream of knowledge would be made instantaneously accessible to all. The web was exclusively intended as a knowledge cornucopia. In fact, Tim Berners-Lee, the inventor of the World Wide Web, literally used the word "utopian" to describe his early vision of it. ([https://qz.com/1568970/tim-berners-lee-on-founding-the-world-wide-web-30-years-ago/]) While this goal of a knowledge machine was partially achieved, with such sites as Wikipedia becoming huge successes and a proliferation of good information sources and the widespread availability of scholarship, little could have internet visionaries known that the explosion of digital information would also spell a plague of misinformation virulency. The modern internet is thus a epistemically hazardous mixed bag of good and bad information with none of the regulation one would expect from an educational institution but all of the knowledge of one, as if a truckload of priceless gems was dumped in a garbage landfill and blended. Moreover, the information deluge does not mean an overflow of knowledge. Information is the precursor of knowledge, but knowledge has a more systematic and organized representational structure. On top of the trouble of having to sort the wheat from the info chaff is the compounding issue of info glut. So all in all the internet is as much an ignorance machine as it is a knowledge machine. To add insult to injury, as Nicolas Carr argues in *The Shallows*, the internet can cause disruptions in attention and focus that degrade the ability to acquire knowledge and think critically. It is not just corrupted information content, but also tainted processes we need to look out for. (By the same token, according to the World Economc Forum nearly half the world population is still offline

([https://www.weforum.org/agenda/2018/05/half-the-world-s-population-is-still-offline-heres-why-that-matters/]). It has to be wondered whether this deprivation of access is more inductive to the net gain in ignorance than all the angiogenic forces of misinformation unleashed by the internet.)

As these examples show ignorance can be created and added to the world. As various technologies such as AI neural networks become better at simulating human actions or producing artifacts with disturbing verisimilitude, such as generated imitations of a person's voice or likeness and even more fully fleshed out audiovisual deep fake videos (all of which as of this writing have been convincingly accomplished), we are entering a brave new world of agniogenesis that could threaten the very foundations of the epistemic trust on which a functioning society rests. So it is all the more important to promote critical thinking skills and a healthy dose of skepticism. Dramatic though the statement may be, it is these skills and habits of the mind which are the last, and yet still most reliable countermeasure against the machinery of deception.

Examining the epistemic form of a lie reveals peculiar knowledge-ignorance interactions. A lie is a message that packages a kernel of ignorance that "infects" or violates the epistemic integrity of its recipient. The liar knows the lie is false and that if believed, it will create a pocket of ignorance in the receiver. A lie is effectual because the victim of the lie unwittingly believes in world that is falsified by it. An effective lie does not lose efficacy because it omits or misrepresents facts in the way that an honest but false statement does. A good lie purposefully recapitulates the surface appearance of the truth and works within the space of believable possibilities while modifying specific details to achieve the goals of the liar. It is planted between the interstices of another's ignorance, becoming a virus-like packet of false knowledge which nests inside the mind of the host, distorting their understanding of the facts.

Traditionally the philosophical analysis of lying is concerned with the definitional problem—what is a lie?— and the ethical problem—why is it wrong to lie? Here, the goal is not to define what a lie is beyond what we have already established in the preceding paragraph as the deliberate manipulation of a

differential in knowledge and ignorance between parties. Nor is it to probe the ethics of the deed. The traditional definition of lying is "a statement made by one who does not believe it with the intention that someone else will believe it" (Isenberg 1973). Yet this definition lacks an appreciation for the epistemic implications and pragmatics of lying. The ethics of lying, furthermore, is extraneous to this additional topic of interest. What is of principal interest is the epistemic and logical form of lies and how these pertain to the investigation of inverted epistemology.

For a lie to succeed a specific, complementary configuration of relative knowledge and ignorance must hold between the communicants. The liar must know that the person they're lying to does not know the truth, they must know that the statement they will utter is false, and they must have confidence that the lie fits their victim's preconceptions such that it bypasses any critical checks that might expose such that recipient is incentivized to reflect on and doubt the message's version of events. Contrariwise, the recipient of the lie must in fact be ignorant of the truth, must believe the statement to be true, and must have a contextual understanding of the events relevant to the lie that is otherwise consistent with and undisturbed by it. Finally the recipient must never come across new information that contradicts the lie. If any of these conditions fail, either the message in question is not a lie or it is be a lie that is found out.

Often such requirements are external to the analytic meaning of the statement itself and reside in the wider social and situational context in which it occurs. An insincere sense of trust must be cultivated in the the recipient by the liar, "buttering them up" to create in them suitable conditions conducive to believing the lie. If the liar has a reputation for lying, successful lying becomes more difficult and the internal sentential structure of even the most ingeniously concocted lie alone will not suffice to gain purchase. A lie that is not doubted, might also cause the victim to refrain from seeking out information that disproves it, but the truth might come out anyway in the normal course of events. Thus a lie unfolds within a backdrop of conditions which determine and constrain it.

An inadequately concealed or constructed lie will eventually come into

conflict with reality, outing it for what it is. An effective lie will however evade this outcome. The fact of that evasion alone says something about the truth, about the possibility structure of the world. For a lie to be believable, it must be plausible, it must map to an approximate location near to the truth within the "possibility shell" of uncertainty shrouding the representation of an event. Within the modal space of a particular event are certain counterfactuals are more adjacent to the objective state of affairs while others are more remote. When someone is framed for a crime, the outlay of the evidence is made to fit them so that the explanation in which they are the culprit is more believable than the alternative. A blatant, bald-faced lie is one that does not fall within the margin of plausibility that is adjacent to a local fact in possibility space. A successful lie works around the truth. You can only frame someone who is framable, who fits certain criteria that plausibly place them at the scene of the crime. To use a trivially excessive example, no one would frame a fictional character or Zeus for a crime, or someone who was obviously incapable of it like a newborn baby or the deceased. A believable lie, by its very believability, works within a margin of flexibility, a band of acceptable deviation from the facts which falls within the overall context of the plausible.

Just as cognitive biases, although methodologically flawed, can still produce a kind of pragmatic advantage or slanted knowledge, so too can lies probe the truth by testing the bounds of believability. A "coverup" is a lie that maps the facts one to one. Any "danglers" or loose ends, would expose the lie. The coverup and the truth share a certain correspondence and symmetry. Similarly, an exaggeration, "a bit of a stretch" amplifies the truth to an excessive degree. It still bears some relationship to the truth. The idiom "to bend the truth" suggests that facts can be permeated with untruths and still retain it likeness. Verisimilitude is ignorance posing as knowledge.

Gullibility implies a deficit in doubtfulness which stems from an inattentiveness to the conditions of uncertainty that shroud events. As stated in Part I, any one belief unfolds within a backdrop of prior belief and is judged in terms of its compatibility with it. What one already believes forms a framework of constraints on new beliefs. Any candidate belief that deviates too widely from

these constraints is flagged for suspicion and earmarked for rejection. Arguably, as Peirce puts forward in the essay *The Fixation of Belief*, the goal of the mental economy is to achieve and maintain the stability and internal consistency of belief. Thinking occurs when doubt threatens to disrupt this placid state of self-assurance. Anticipating cognitive biases, Peirce says that the objective is not to obtain the truth but to resolve doubt such that the stability of prior belief is restored. The boundaries defined belief constraints define the bounds of what is considered plausible.

In the absence of external information, an agent's subjective "belief bubble" is a closed system with finite inference value about the world at large. This ring of beliefs must occasionally open its doors to let in new information and add additional beliefs to the repository in order to accomodate new developments and scale with time. This opening is selective however. The constraint framework's network of interdependencies generally leaves open a "clearing" for potential candidate beliefs that it would accept and the rest are to be excluded. An effective lie threads the needle of this opening, finding a place among the other beliefs by posing as admissible when it would otherwise be locked out of the system. The structure of belief in other words fixes it so that there are distinct receptivities and vulnerabilities particular to each individual, or what in cybersecurity is called a threat profile, that uniquely expose them to exploitation. It's not enough to say the vicim of the lie does not know the truth. Their entire system of belief does not protect them from falling for the deception.

This account thus far has made it seem that the liar enjoys a significant information advantage over their target. Not so. Much as ignorance makes the recipient susceptible to trickery, the liar is also in danger of their own ignorance and cognitive limitations. Keeping track of their own deceptions incurs a kind of "maintenance cost" or memory attrition. They have to keep two sets of books, making sure they keep their story straight and never confuse the lie with reality —which continues unfolding along its course without respect for the lie—and risk forgetting their own spin. They must remember what is true and what is falsely believed to be true and be sure the two never mix. As Dennett (Brainstorms) points out, lying requires tracking a bifurcated model of events.

Cognitive effort must be expended making sure the two parallel streams do not come into conflict. Lying prompts a build up of logical dependencies between statements as the liar weaves a web of fictions that must cumulatively pass bullshit checks which become exponentially harder to pass as the web of lies grows in complexity. A liar may be forced to fabricate secondary "supporting" lies, alibis and counterpoints to correct for narrative inconsistencies and rebuke suspicions. These compounding requirements become burdensome to sustain and overload memory. When lying there is always the risk of forgetting that one has lied and to intuitively and naturally re-pattern one's understanding back onto the truth and to thus act in a way that contradicts one's initial position. In contrast a true and honest statement has the world to back it up, and it is perhaps for this reason that honestly and "coming clean" is experienced as a relief. No more effort must be expended keeping up appearances and mental states can relax back into realignment with the single stream of manifest reality—or at least, what we earnestly believe to be the case.

Both communicants in an episode of lying are therefore subject to to various manifestations of ignorance and neither is immune from its effects on their thinking. While it might be intuitive to assume that the liar has a distinct knowledge advantage, at least in the narrow sense of their capitalization of a fact of which the other is ignorant and which seeds the content of the lie, even the liar is ultimately the casualty of information entropy. The information cost involved in sustaining the lie against the truth becomes overwhelming unless a substantial effort is exerted to cover up and counter reality. From a purely cognitive perspective, is much harder to pull off a lie than to tell the truth. To tell a lie is to enter into a war of attrition with reality; while to tell the truth is too gently side with it. In the immortal words of Shakespeare:

> Nay, indeed, if you had your eyes, you might fail of
>
> the knowing me: it is a wise father that knows his
>
> own child. Well, old man, I will tell you news of

your son: give me your blessing: truth will come

to light; murder cannot be hid long; a man's son

may, but at the length truth will out

When Galileo faced trial for for his heliocentric views more was at stake than an interesting if practically inconsequential astronomical discovery. If that were the case the behavior of the ecclesiastical authorities would be nothing short of petty. Quite the contrary. The question of the position of the sun in real space was also the question of the position of the sun in mental space, a position which pinned in place both the celestial firmament and the ideological firmament of the Church's teaching. For the prevailing view of the world at the time to hold in all its dependencies and relations, the geocentric view had to be upheld, the authorities feared, or else the elaborate machinery of an entire belief system that glued together the culture of Christendom would risk slipping out of joint and collapsing into tatters.

Irrational and anti-intellectual as though the Church's treatment of Galileo may seem from modern sensibilities, from the standpoint of its fundamental cultural role in the 17th century it makes a kind of sense. As much as the clergy's actions reeks of ignorance from one angle, from another there is an inkling of sober wisdom. "People can't stand too much reality," Swiss psychoanalyst Carl Jung once said. In the dark world of early modern Europe, where plagues were frequent wretched occurrences, infant mortality high, the pains of living ever-present and the average lifespan short, it was encouraging to believe that humanity is at the center of the cosmic drama and anxiety provoking to dwell on the prospect of human peripherality. It takes a certain kind of mental fortitude and stoicism to accept that humanity might be a cosmic afterthought, especially when life is hard. Indeed religiosity and comfortable living standards tend to negatively correlate: poorer countries tend to be more religious and richer

countries less so. ([https://advances.sciencemag.org/content/4/7/eaar8680.full]

If the world's religions have one thing in common it's that they almost universally dignify human existence, giving meaning to suffering and solace in misery. To the extent that religion and science are antagonistic it could be because the alienating and inhuman reality revealed by science threatens what comfort and moral order those who are religious and suffering have. Yet antagonism more often than not signifies mutual incomprehension, mutual ignorance. To be truly at odds two rival ideologies must deny any truth and any possibility of being true to the opponent. Reconciliation means the realization that the truth is bigger than what either opposite holds, an intolerable outcome for belief systems defined by a rigid doctrine that claims to profess the final word.

Truth is at times described as enabling, enlightening, and freeing, at other times ugly, fatalistic, and inescapable. There are truths that will set you free or harsh realities, truths that inspire and vindicate or truths that demoralize and crush. We speak of the hard facts, insensitive, neutral, and impassive. Most of all however, truth is productive. Truth is the light which reveals the hidden. Like wealth, it's generally always better to have more truth than less and to be with too little of it is to be impoverished. To be in possession of the truth is to be able to move forward, to be aligned with reality. Truth is essentially constructive. knowledge is power.

Yet from a different perspective truth is destructive. When the truth touches and transforms a cluster or ignorance, it's like an acid dissolving through foam. It eats through comforting preconceptions. Truth can be destabilizing. Because ignorance can have positive reality, can cause events and organize society, to dispel ignorance is sometimes disruptive. The insistence that more truth is always better and to push for truth at all costs ought to be reevaluated in light of the pragmatic consequences of that truth rather than in terms of its intrinsic value.

The idiom "what he doesn't know can't hurt him" intimates that some knowledge can disorient. Suppose that an astronomer viewing the night sky discovers that a massive asteroid like the one that drove the dinosaurs extinct is hurtling toward the earth. Its impact will wipe out humanity. Nothing can be

done to redirect it, and iron clad calculations show there is only a 50/50 chance that it will narrowly evade impact. Should this astronomer be obligated to tell the world about these findings and by doing so spoil what time humanity has left? The result would be the widespread outbreak of anxiety, helplessness and despair. What's to be gained from withholding this knowledge is instead at least the undisturbed illusion of normalcy until the very last moments. In the event that the asteroid misses the earth, furthermore, it would be like nothing ever happened, in which case there is no gain at all to be had informing the public and causing everyone needless stress.

In *The Grand Inquisitor* by Russian novelist Fyodor Dostoevsky, Christ returns to earth and visits the city of Seville in late medieval Spain at the height of the inquisition, during which hundreds of accused heretics were persecuted and burnt at the stake. Upon arrival Jesus performs various miracles and wins the adoration of the townsfolk. "People are drawn to him by an invincible force, they flock to him, surround him, follow him. He passes silently among them with a quiet smile of infinite compassion. The sun of love shines in his heart, rays of Light, Enlightenment, and Power stream from his eyes and, pouring over the people, shake their hearts with responding love. He stretches forth his hands to them, blesses them, and from the touch of him, even only of his garments, comes a healing power." He restores sight to a blind man and resurrects a recently deceased seven year old girl from her unburied coffin. Just then the Grand Inquisitor, the chief overseer of the inquisition's operations, strides over to where a crowd has gathered to ecstatically celebrate Jesus's return. Without hesitation he points a finger and orders Christ's arrest. The crowd parts with striking obedience and allow the guards to seize Jesus and then bows in abject deference to the Grand Inquisitor as they take him away.

Christ is locked in a cell and later that evening the Grand Inquisitor enters. Rather than interrogating Christ, who sits in meditative silence throughout the ordeal, he engages in a long theological and philosophical monologue explaining his reasons for committing the apparent ultimate heresy and contradiction of imprisoning the central figure of the faith he claims to champion.

The inquisitor asks Christ why he has come to "interfere with us." He

explains that Christ's promise to free humanity fifteen centuries prior had resulted in the scenes of wretchedness and despondency he saw on the streets of Seville. "Man was made a rebel; can rebels be happy?" He explains that by giving the Keys of Heaven to Saint Peter, Christ relinquished spiritual authority to the Church, and so it is now their responsibility to govern the moral direction of mankind. He refers to the Temptations of Christ, in which Jesus, exiled to the desert for 40 days and 40 nights, engages in spiritual battle with Satan who attempts to ply him with three tempting offers for his soul. The inquisitor identifies with Satan, whom he does not refer to by name but instead calls a "high spirit" of incredible "power and intelligence", and who he sees as the voice of earthly reason and pragmatism. Knowing that Christ is famished from exile in the wilderness, Satan entices him to turn stones into bread. Christ refuses, stating that "man cannot live on bread alone", implying that spiritual discipline is more important than corporeal satiation. The inquisitor scoffs at this. "Do you know that centuries will pass and mankind will proclaim with the mouth of its wisdom and science that there is no crime, and therefore no sin, but only hungry men?" he says. "'Feed them first, then ask virtue of them!'—that is what they will write on the banner they raise against you, and by which your temple will be destroyed." The official goes on to say that it is an impossible expectation to hold the masses to Christ's spiritual standards. For them freedom and knowledge of higher truths is a maddening burden, and so it is the duty of a paternalistic authority to hide the truth from its followers. It is essentially immoral, he argues, to leave the mass of humanity wallowing in the terrible dilemmas of freedom of choice. The inquisitor anticipates Jean-Paul Sartre's existentialism and his claim that we are "condemned to be free", that the question of what to do with ourselves in this world is overwhelming and there is always the temptation to live inauthentically by deferring to authorities.

The point of mentioning Dostoevsky's profound parable is not to single out or even draw attention to religion or theology, but rather to use it as an allegorical window into the complexities of truth and ignorance. The paternalistic institution capitalizes on hidden truth, for fear of the chaos it would inflict should it escape the confines in which they have bottled it, like deadly

radioactive material leaking from an opened lead canister. It commits what Lindsey McGoey in her book *The Unknowers* calls strategic ignorance but with a slight twist. By strategically withholding the harsh truth about human nature, the Church allows humanity to live in childish innocence, utilizing this discrepancy in knowledge to uphold civilization rather than let existential dread consume society.

In today's complex world, there are numerous threats that defy solution. Climate change could radically alter the world for the worse, nuclear war is an ever-present possibility, mass extinction claims the lives of ever more species, water scarcity afflicts millions and is only getting more severe. It would be insanity to simply bury our heads in the sand and pretend these problems don't exist. At the same time, stressing the genuine gravity of these crises is probably too much for many to bear.

In 2019, OpenAI, a research organization that develops artificial intelligence, disclosed to the public GTP-2, a deep learning language model that aside from some occasional semantic kinks and nonsensical constructions generates strikingly humanlike text. Fearful of how this technology could be abused yet wanting also to foster more understanding, the OpenAI team released a limited version of the software to the public, with the caveat that they would withhold the complete training set the AI used to learn its full capabilities. After some experts made the case that GTP-2 could also be used to detect fake news rather than just create it, OpenAI later reconsidered their position and released the full version. However, the organization's initial hesitancy speaks to a genuine concern. Is there some knowledge the mass of humanity is not ready to handle?

"The most merciful thing in the world, I think, is the inability of the human mind to correlate all its contents. We live on a placid island of ignorance in the midst of black seas of infinity, and it was not meant that we should voyage far", wrote H.P Lovecraft. One wonders if there are indeed things that humanity is not meant or not yet ready to know. In the science fiction franchise Star Trek, the Federation's Prime Directive, its most important rule, is to never interfere with the natural development of civilizations that have not yet realized interstellar

spaceflight. The reasoning is that such meddling even if well intentioned could radically destabilize those civilizations. Introducing technology and knowledge too advanced for those civilizations could yield unintended effects and disturb the culture's natural state.

In spite of all its benefits, the radical scientific pursuit of truth and knowledge also brought us nuclear bombs and other dangerous insights into how to craft biological and chemical weapons. Science fiction and now increasingly genuine scholarship speculates whether artificial intelligence may one day become self-aware and view humanity as a threat, turning our technology against us and destroying society. The knowledge of fossil fuels that powered the industrial revolution also gave us the global warming which looms ominously over the 21st century. Automation could upend the economy, spur mass unemployment and lead to incalculable societal volatility. Like Prometheus we play with fire, dabbling in forces beyond our full comprehension. The so called law of unintended consequences, which speaks to the idea that even presently well informed actions can lead to unforeseeable results, ought to inspire humility. Examples such as the Hindenburg disaster in which an idea that sounded good on paper ended in fiery calamity abound. Like opening Pandora's box, once unleashed such epistemological forces are irreversible and unpredictable, paradoxically increasing uncertainty and ignorance by creating new and unfamiliar contingencies. There is hubris in the belief that we will always be the masters of our knowledge and shrewd enough to control it wisely.

This discussion about the destructive power of truth inverts the typical tendency to view knowledge as a pure good and ignorance as a pure evil. The reality is much more nuanced. Only by accounting for the complete epistemological picture that appreciates the complex interplay of knowledge and ignorance can these complications be properly understood.

Perhaps you might be left with a feeling of pessimism after having considered all the arguments presented in this book. One of the most tragic consequences of the realization of the true scope of ignorance is a sense of the paltriness of one's own knowledge. To grasp ignorance in its full majesty is to

dispel the Dunning-Kruger effect, the comfortable delusion that we're more knowledgeable than we really are. The surface is only ever so much as scratched. So many books will go unread, so many experiences will never be. Wondrous worlds will remain undiscovered and unexplored.

This mystification before the unknown speaks more to the human condition, however, the sting of mortality, than it does any theoretical epistemic limitation. Yet in the end they are deeply related. No matter how fully we live our lives or how avariciously we pursue knowledge, it seems as though we will only ever chisel a pebble off of Mount Everest . And the more we are aware of how much is out there yet to be discovered, the more this wistful dissatisfaction sinks in.

The modern internet phrase FOMO (fear of missing out) attests to this uneasy discrepancy between our ability to discover what's going on and our incapacity to experience all of it. The internet itself displays the paradoxical tendency for an spectacular increase of information to have the opposite effect of informativeness, to say nothing about it actively increasing ignorance through rampant misinformation spread. As Nicholas Carr argues in *The Shallows: What the Internet is Doing to Our Brains* , the internet's overabundance can interfere with concentration and its plethora of competing information signals lead to pervasive opportunities for distraction.

Our very tools of discovery and enhanced awareness painfully enlarge the recognition of our own obliviousness and cognitive finitude. An avid reader gets the same sensation when they look over all the books they've collected and realize they've only read a fraction, or the film aficionado when they realize they've barely made a dent in the classics. DaVinci for all his gifts knows less about science than a bright high school graduate. Even the most enviable polymath, the most knowledgeable person in the world, even the most knowledgeable person *conceivable* can never be a true know it all. According to professor of psychology Paul Reber of Northwestern University, the human brain can encode about 2.5 petabytes of information, roughly equivalent to a million gigabytes. That may seem like a lot, and it is. But there is a lot more out there ([https://www.scientificamerican.com/article/what-is-the-memory-capacity/]. The visible universe, contains 10^100 bits of entropy, and that's just

the part that's visible ([https://www.scientificamerican.com/article/information-in-the-holographic-univ/]). That is a difference between a 25 with 15 zeros after it and a 10 with 100 zeros after it. Admittedly, information bits are not the same as knowledge, any more than a yes or no answer without a context is. Information gives knowledge something to work with; they are not substitutable. Knowledge tells you what to do with information, but knowledge without information is like eyesight without light. In any case, the raw information value of the physical world means on a basic level that mere mortals will never get abreast of it all.

It becomes even more dizzying when one considers mathematical infinity and what that means for one's limited knowledge by acquaintance. Presumably it would take an infinite memory to know and remember each of the infinite digits of the decimal expansion of pi. In the broadest strokes, the difference between the information capacity of the brain and the universe shows mathematically that there is more out there than you could even fit in your head, let alone become acquainted with to know in a lifetime. Or even for that matter, a quindecillion lifetimes.

Of course nobody is interested in gaining an intimate knowledge of every grain of sand, every particle of dust, every leaf in the wind, blade of grass, or space rock hurtling through the intergalactic void. The most powerful and robust knowledge encodes general principles. Science folds the redundancies and similarities of phenomena into compact schemas and interprets the token in terms of the type. Many objects we don't even have names for, so generic and trivial is their existence. Imagine if we bothered to name every blade of grass. So much goes unnamed and, according to the logic of language, therefore undiscriminated and undistinguished as an individual token. Indeed it is just this sort of economization that compiles indistinct elements into categories that allows us to overcome our systematic ignorance of the particulars. (Although as in the case of social and racial stereotypes, this mode of general thinking also can misfire and spawn its own forms of ignorance. It is eminently possible to overgeneralize.) It may not be easy, and in the end we might trick ourselves into believing we truly understand it all when we don't quite, but to be able to affirm

valid generalizations is a huge improvement over being confined to ignorance about each instance of phenomena rather than the class as a whole.

In moments of uncertainty or intellectual challenge when the depth of one's ignorance is painfully clear, you experiences what might be called *ignorance overload*. For all our powers of understanding and now matter how learned one is, there is always a *more-than*, a clamoring, incalculable totality beyond one's scope only a fraction of which any one individual can hope to be acquainted with in their lifetimes. It is this pressured totality of missing knowledge which gives thrownness its sharp edge, for being thrown into the world is problematic because beings are surrounded by ambiguity.

Again much, even most of this ignorance stuff is not worth knowing, or even better off not knowing. Yet not only is there this vastness of the outer unknown, but there's also the vastness of an inner unknown. Secondary to the problem of the overabundance of the unknown external world, is the overabundance of ways of looking at it: the profusion of different interpretations, modes of acquaintance, theories, allegations, conjectures, stories and speculations; and on even *on top of that* the additional problem of how to select which of these is right. Then there is also the vast inner negative space of unthought thoughts, inconceivable conceptions, unarticulated explanations, to complement this inner space of selective possibilities. Ignorance seeps in everywhere, not only in the outer world, but in and between our very modes of acquainting ourselves with it and deciding what to believe is true. To the clutter of the external world is appended the internal clutter of the mental world. It's not merely that things out there are overloaded, but so is, or at least so can be the mind.

The mental world however, eventually comes to a resting state of habit which borders on the senescent. One cannot keep fighting ambiguity forever. All thinking and doing is ultimately *pressured*, by time and the facts at hand, by the limitations of memory and one's powers of conception or capabilities of action; by the need to decide and for one eventuality to come to pass rather than any other alternative. And like anything that is pressured, it either hardens or it breaks. What's jumbled must be straightened out, if there is any semblance of rationality to our thinking. Possibilities sharpen into actualities. The criteria for

what's possible is less strict than for what's actual, which is to say, there are more possibilities than actualities. The former collapses into the latter. Eventually, the two become the same. Like the wave function collapse, uncertainty must eventually resolve into certitude. The amorphousness which contains all possible shapes in its shapelessness becomes one. Fuzzy foresight in time yields to clear hindsight.

Even writing this sentence there is a subjective sensation of tumult and inner competition between alternatives. Of the many things I could write I must settle on one in the course of deciding my thoughts. There is an opportunity cost to all knowledge or proactive activity. All must die so that one may live. I will never think those abandoned thoughts; venture down those paths. They were nipped in the bud. Thinking therefore is invariably a process of creative destruction. I can't help but feel that in the process of selecting one option I am correspondingly eradicating another. I can think of what is excluded as the price I pay for selecting the one option, the information cost of deciding on it. In sum, there is always more to be dealt with, more to consider than can be considered. This pressure forces what will ultimately come to pass into being.

On a bright note, the vastness of the unexplored is also exhilarating and energizing. There is work to be done. Everyone has a unique experience, their "word-line" that traces out a path through the universe will expose them to a facet of reality no one else, no matter how intimate with them, will ever exactly share. The fact there is so much we don't know keeps the life of the mind from getting too boring. Ignorance overload is also an abundance of richness. Ignorance *scarcity* would strangely enough almost be impoverishing, because it implies that one's environment is numbingly redundant and uninteresting.

The Psychology of Ignorance

While ignorance has a certain mind-independent reality to it, or at least can be fruitfully characterized as such, undoubtedly it plays an extensive role in human psychology. It might be more intuitive to place this psychology at the heart of the analysis of ignorance, but to do so is to give into the temptations of

anthropocentrism. In my view, the psychology of ignorance while fascinating and important is peripheral to the deep informational structure of it. Be that as it may any proper treatment of the subject deserves an appreciation of the intricate ways in which ignorance weaves through mental life.

How to make sense of the problem of consciousness in information theoretic terms? In previous chapters we touched on the burgeoning field of cognitive thermodynamics and the theory that consciousness is associated with the maximum entropy state in which the maximum of brain networks are synchronized ([https://www.frontiersin.org/articles/10.3389/fnins.2019.01106/full]). Recalling Jayne's principle of maximum entropy, brain states which have the maximum amount of parameters have the maximum number of available options for interconnecting and exchanging information The more entropy, the more possibility, the more information, the more consciousness. Conscious awareness therefore is connected to how many possible configurations different neural networks can enter into with each other, which maximizes the information flow between subsystems. The more free moves a network has to connect with others, the more this correlates with elevated levels of conscious awareness.

These correlations between the amount of available information in the brain and awareness leads to interesting implications when viewed with respect to the premises of this work. According to the ideas presented here, knowledge is associated with low entropy —ordered—states. knowledge is closer to necessity than possibility, we know something when the case cannot be otherwise than what we know it to be. We typically associate consciousness with knowledge-as-awareness. To be conscious is to be aware and to be aware is a primitive form of knowledge. This leads to a conceptual mismatch: if consciousness is associated with high entropy and therefore ignorance, how come we associate awareness with knowledge, which is a low entropy, ordered state?

The idea that consciousness correlates with maximal entropy would mean that consciousness is a state of ignorance . However, ignorance as pointed out in chapter 7 is the epistemic correlate of possibility. To know is to know with necessity that such is the case while to not know is to tacitly affirm that all

potential states are equiprobable. Possibility entails that there are more parameters in play than strict necessity . In this context it consciousness is related to ignorance as potential knowledge (awareness) rather than as the absence of it . Ignorance, in other words, can also mean a kind of slack or openness, the possibility for an event to swing one way or another. Thus, a set of networks composing the brain which has the maximum number of possible configurations opens itself up to the world by allowing itself to be receptive to alternatives. This chimes with previous commentary that observed how it is not what one knows that counts, but what one can knows. The mind which starts slow but rapidly gains speed as it learns will soon exponentially out pace a mind which starts with a fixed sum of knowledge, even if this sum is initially much higher.

Recall the idea of senescence as a kind of freezing over of fluid cognitive flexibility. Senescence is what happens when knowledge grows stale, when, the leading, growing edge of learning sputters out and slows down as the mind's capacity to update itself on the basis of new information locks up. knowledge tends towards metastable stasis over time, ultimately, what the organism comes to know is catered to accomodate to its needs rather than to objectively represent the world as it is in itself. But if knowledge reaches pure stasis, that is, if incrementing it and growing it over time becomes more difficult because the novelty learning requires becomes rarer, knowledge representations are at risk for going outdated. knowledge that fails to continually revolutionize itself terminates , somewhat counterintuitively, in ignorance. It ends in fossilized representations of a world that no longer exists rather than the latest conditions. knowledge has a certain lifecycle that must be respected if it is to be preserved. Brains function optimally when they are in metastable states, which is to say, when they are sensitive to minor perturbations but nevertheless reconfigurable and reflect those changes in its plastically adaptive neural network. (Fingelkurts, A. A., Fingelkurts, A. A., Bagnato, S., Boccagni, C., and Galardi, G. (2014). Do we need a theory-based assessment of consciousness in the field of disorders of consciousness? Front. Hum. Neurosci. 8:402. doi: 10.3389/fnhum.2014.00402). As such, the brain is neither a perfectly ordered nor a chaotic system, it straddles the middle ground where interesting stuff happens.

knowledge crystallization can be associated with cellular compartmentalization, the process by which hydrophobic lipids self organize into the layers of cell membranes. "In the brain a compartment can be thought of as a number of cell networks that are synchronous, or functionally connected." ([https://www.frontiersin.org/articles/10.3389/fnins.2019.01106/full#h8]). Compartmentalization, as the formation of spontaneous order, signifies a reduction of entropy rather than an increase. Therefore it is entirely consistent that knowledge is a metastable low entropy ordered state while consciousness is a diffuse high entropy state. Awareness, then, is the possibility for knowledge, while knowledge is a actualized, collapsed, crystallized substate that is contained in that space of possibility.

knowledge and awareness differ in that the former displays propositional structure while awareness is sub-symbolic. knowledge is *re*presentational and awareness is *pres*entational. Propositional knowledge describes something of which a subject is aware, while awareness is direct and unmediated. Even though they both differ, they also complement each other. One can only come to know about something that one is aware of. In addition, knowledge is about deploying content while consciousness is about access to that content. To utilize knowledge we must be aware of it. The awareness of that knowledge is not another redundant representation of a representation but an immediate, given *presentation* of it. Think of how inefficient it would be if all knowledge representations needed another one simply to confirm that we knew them. Another would be needed for the secondary representation, and so on ad infinitum. More straightforward is to free up access to all available states without having to rely on any encoding mechanism to reference stored knowledge. So it's more accurate to say metaknowledge is being aware we know rather than knowing that we know. Consciousness is the *possibility of access* to local knowledge states. Consciousness is therefore a state in which the maximum available micro-states—which by mathematical necessity must each have lower entropy than the macro-state—are made accessible.

If this account is true it would represent a profound example of structural ignorance. Consciousness, which ironically to have such a hard time

understanding the true nature of, is really a tissue of ignorance as described by low level information. It is ignorance in the sense of openness, tentativeness, degrees of freedom, unresolved complex probabilistic activity. It is not closeness, decidedness, strict and simple deterministic behavior. What you have in this second case is not consciousness but habit, reflex, behavior so redundant that there is no reason to need to notice it.

One can speculate from these ideas about various mysteries of consciousness such as hallucinations, dreams, or psychadelic experiences. Normative waking consciousness is typified by homeostasis. It regulates an ordered gestalt of perceptual elements which track and correspond to more or less faithfully the correlated changes of the sensory receptor inputs as shifting arrays of stimuli stream into them over time. In dreams or hallucinations or psychedelic episodes, the coordinated coupling between consciousness states and sensory input streams is interrupted. The brain therefore loses the cues by which it can associate inputs with corresponding local representations, and so the mind runs off its tether. Like a guitarist who stops playing the repetitive hook of a song and goes on a quirky solo, the ensuing sequence of conscious states is irregular and unique. In sleep, the lack of sensory input is what triggers these uncoordinated cascades. Psychoactive drugs, however, prompt an action which overrides the sensory stimulus and overwrites it. Similar explanations from the predictive processing viewpoint have been put forward ([https://aeon.co/essays/use-uncertainty-to-leverage-the-power-of-your-predictive-brain]). Essentially such substances break the loop and incoming sensory signals fall out of sync with outgoing prediction signals. The result is the experience of bizarre perceptions as the unlinked prediction signals cease tracking sensory inputs and instead trigger without constraints.

> On the predictive processing view, predictions are essentially compressive, where higher levels of the hierarchy dealing with more abstract or invariant features of reality, and tracking longer timescales – strip the sensory signal of redundancies: that is, anything that's not relevant to adaptive action. Under psychedelics, the 'reducing valve' is opened. From a predictive processing perspective, these drugs can be understood as loosening the grip of engrained

and rigid expectations about how sensory signals are caused, allowing the brain to create novel hypotheses about the world and how we relate to it.

The deep questions of consciousness could continue to be probed by the concepts we now have at our disposal.

Common sense tends to differentiate emotion and cognition as two separate categories of mind. Emotion is irrational and impulsive, while thought is rational and deliberate. This partition however, is empirically unsupported. Emotion and cognition are not sharply distinct. They interoperate within in the same functional networks and it is valid to speak of emotional cognition or cognitive emotions ([http://174.143.211.73/\~ellenlau/courses/nacs642/Pessoa_2008.pdf]). If cognition is, as I argue, fundamentally about uncertainty reduction, of sorting the actual from the possible, and if emotion is woven into cognition, then naturally it might be inferred that emotions are related to negentropic processes in some important way. The interactivity of emotion and cognition has led some scholars to speak of hot versus cold cognition, which designate motivated reasoning on the one hand and purely calculative logical processing on the other. (Brand, A. G. (1985–1986), "Hot cognition: Emotions and writing behavior", JAC, 6: 5–15, JSTOR 20865583)

It's easy to be skeptical that emotion has the same negentropic role as cognition and it's hard to see how it can assist in its execution. After all, what can be more fraught and perplexing than one's exasperating feelings, the overwhelming complications of love, the dread of death, the stifling frustration of a bad mood? Emotions often make things messier rather than simpler. While the cognitive in contrast is seen as reserved and composed and "cold", emotions are boisterous and noisy and "hot". Though both may work together some of the time, they are an odd couple and often enter into loggerheads. We speak of "losing one's cool", of slipping into sloppy and emotionally clouded thinking, ceding rationality to impulse. We don't find prudence in someone who is "hotheaded", temperamental and prone to outbursts. Yet this puts things the wrong way around. Cognition serves the emotions. As Schopenhauer framed it, the intellect is subordinate to the will and is tasked with delivering on its desires.

Except that Schopenhauer repeats the mistake of treating them as separate and distinct. Really, in the fully functioning animal, thinking and feeling are part of the same cohesive system. So it's natural to ask: what exactly is the cognitive function of emotion?

Various theories have been proposed (ibid). Such as that emotion is a means of processing to rewards and punishments, or that it is the engine that drives motivation. One that I favor here however is that emotions have an *evaluative* role; they are involved in the cognitive appraisal of the meaning of a given event and are part of fast, automatic system I unconscious processing (Rolls, E. T. Emotion explained (Oxford University Press, Oxford, 2005). Emotion is implicated in *action readiness* it prepares the agent to act. Emotion, which is rooted in evolutionarily conserved, primitive subcortical brain structures, emit quick gut reactions to stimuli, triggering with a firm and unmistakable tone that sub-symbolically and pre-consciously communicates to the relevance of an event to an organism. Emotions convey information without having to wait for discursive, language-mediated cerebral cognition to slowly and consciously work through the details. Emotion bypasses the guesswork and gets things moving without having to hover in a state of pending deliberation that forces one to stop and consider the relative merits of various alternatives before acting on a decision.

Emotion is then appropriate for action at short to mid-term timescales, providing a quick précis or snapshot of an event's visceral meaning not to guide action but prompt it. Purely deliberative cognition, however, is more appropriate for mid-to-long term decisions where there is enough distance to consider the possibilities. As one might expect by looking at the stacked layers of the brain, the structures regulating emotional decision making stands above the instant, low level preconscious reflex emitters controlled at the basal level of the brain stem but below the cerebral cortex. They are in the middle, embodied by the limbic midbrain centers that reside beneath the long-term higher executive functions at the top of the hierarchy in the cortex. Quite literally, emotions are situated at the center of mental life.

What about destructive, misplaced emotions? Emotions that lack wisdom,

which misguide and sow ignorance? Blinding rage or debilitating depression? The perennial cause of tragedy, such volatile emotions indeed create rather than solve problems. For the depressed person this is all that is true of it. Such emotions are no different in terms of their cognitive value than false beliefs, positive ignorance. While true emotions exhibit knowledge, emotional intelligence, false emotions like blind rage or excessive guilt, are misbeliefs. (Arguably, however, depression is not an emotion, though sorrow is. It is an emotional disorder and therefore does not represent normative functioning.)

When one feels something there can be no doubt that one feels it, although when one thinks something, that thinking might be composed entirely of doubts. Emotions are a clear signal. Unless severely disordered, few people have trouble putting a label to their current mood. Emotions express a clear, bold signal. Phenomenologically they have a definite, sure character, they assert themselves over us. Once an emotion has announced itself, it is difficult to consciously control and suppress it until it has run it is course. It dictates the priorities about what is then thought about. The actions which emotions beg for must be completed for it to abate and withdraw. The forwardness of emotion cuts through uncertainty. We do not have to decide what we feel, it decides for us. Unconscious emotions make up our mind for us. Thinking in contrast, is a state of uncertainty as to what to decide, hence its deliberativeness. This way of looking at emotion is intuitive because it does indeed seem as though emotions are something which happen to us like the weather. If that weren't the case, nobody would ever feel sad, because why would they choose to?

Emotions can be split into rational and irrational subdivisions. It is rational to fear a starving tiger or raging bull, but irrational to fear that a tree branch scratching against the window is a ghost. Here, however, the connection between emotional states and cognitive beliefs comes into view. If one has reason to believe the tiger or bull is an imminent threat, then it is rational to fear it. If they are safely behind a partition, this fear is unnecessary. What fear, or any other powerful emotion does tell us, is that there is something which must be addressed in an immediate and decisive way. It does not articulate or conceptualize the true import of this potential threat, but if we were emotionless

robots, it would take too long to process it in a discursive purely cognitive manner.

A powerful emotion in a sense binds the emotive agent to a behavioral contract, telling them that why must do what their heart asks of them or they will continue to feel unfulfilled. It tells them to take action in a visceral, commanding way. The connection between emotion and the bigger picture of cognitive information entropy reduction is therefore bound up with a strong emotion's unambiguous and confident character--a sudden mood is a definite spike of motivation that cuts through swaths of indecision and provides a clarity of purpose. Emotions therefore reduce uncertainty by clearing away any ambiguity about what matters to us. It appraises the basic significance of an object for the emotive agent with respect to their needs and goals.

Faith is one of the most peculiar of human emotions. It bears the form of a belief in a truth, but it defies the rules of rationality. Simultaneously baffling and inspiring, faith presents an interesting edge case in the study of the emotionality of ignorance.

Is faith ignorant? For many who are deeply religious this insinuation borders on the insulting. Two things must be said. One, ignorance is not always a bad thing, there are constructive variants. Secondly, by any and all epistemic criteria it is hard to see how faith passes the test as knowledge. Almost by definition to have faith is to not have arguments or evidence for why you believe it, and if you did it would be knowledge. At the same time we tend to put someone with true, pure faith above someone who is merely stubborn or outright delusional. Such pure faith differs from fanaticism, perhaps by degree or perhaps in kind. It is measured but unyielding, totally clear and unclouded in its total conviction and certainty yet upstanding and dignified. Where does such a willed sense of absolute truth come from?

In my view faith is a complicated mixed epistemic product of both knowledge and ignorance. In essence it is a response to epistemic incompleteness, human finitude. It is the acceptance of an irreducible intuitive recognition that the world has a purposeful order to it, and after Leibniz, that this is a sign of a higher

power working through it. Faith is therefore necessary or at least pragmatic in a life pervaded by unanswerable mysteries and undecidable ambiguities. How we should feel is never a simple logical decision, there are no established rules for how one should decide what is right or wrong, or which emotions are valid. In some critical circumstances the requirements for the rational justification of an action are unfulfillable. One is simply undersupplied with the necessary information and the moment will pass before you get the chance to do anything, assuming it were even possible to acquire more facts. In moments requiring such decisive action, is it better to step back and doubt, risking inaction, or to push forward trusting in one's basic competences and in what good there is in the world to show through for you? This question is itself undecidable, for how could anyone ever answer before the fact whether action or inaction in an information imperfect situation was destined to be the right choice?

Choosing doubt and retreating to await the acquisition of better information may never pay off. Such a person arguably would be on a worse footing than someone who chose to act such that they at least altered the world according to their actions and thereby made it more predictable given that parts of their environment then fall under the influence of their agency. They change the world and therefore gain information about it. The overconfident, in contrast go to far in overestimating there ability to conform the world to their will. They are vainglorious, and wade into a situation in which they are in over their heads. In moments of undecidable uncertainty, confidence can either win the day or lead to worse outcomes than simple inaction or passive observation. Both behavioral policies, and their various outcomes, are entirely contingent on the specific circumstances in play. Both action and inaction can result in wins or loses. One thing is clear however, in the net, action must always lead the way because one who committed to a policy of always relenting would never get anything done. It is those who take action, despite all the mistakes they make doing so, who ultimately lead the way. They may do so out of emotional need, or because they must firm up their conviction to act decisively in the middle of a muddle. Faith allows one to boldly act in situations where there are rational criteria to do so.

Faith, though it may sometimes taper off into wishful thinking or denial,

therefore is an act of boldness before the objectively uncertain and a willful choice to embrace an inner feeling of certainty despite it. The faithful claim to know what they have faith in, but in reality they do not know but they seek to adopt the psychological disposition of someone who knows in a situation where it is infeasible to do so actually. Faith conjures certainty, which alleviates the irritation of uncertainty and allows one to act with a clear head.

The faithful have more than simple confidence, they believe that the ultimate causes of the universe are with them. They have what James called the Will to Believe. In *The Varieties of Religious Experience*, James' acclaimed study of religious psychology, James speaks of the "reality of the unseen." He criticizes the rationalistic, scientific mindset to discount and belittle certainties and convictions arrived at through an emotive process of faith. To properly understand the psychology of belief, the fact that mental abstractions and ideals have actual concrete value to those who believe them must be appreciated despite these mental objects having not been reached through a strictly logical process. Indeed, James asserts, conclusions arrived at by dry reasoning often have less sway on people than those that are arrived at by organic faith:

> … we have to confess that the part of [mental life] of which rationalism can give an account is relatively superficial. It is the part that has the *prestige* undoubtedly, for it has the loquacity, it can challenge you for proofs, and chop logic, and put you down with words. But it will fail to convince or convert you all the same, if your dumb intuitions are opposed to its conclusions. If you have intuitions at all, they come from a deeper level of your nature than the loquacious level which rationalism inhabits. Your whole subconscious life, your impulses, your faiths, your needs, your divinations, have prepared the premises, of which your consciousness now feels the weight of the result; and something in you absolutely *knows* that that result must be truer than any logic-chopping rationalistic talk, however clever, that may contradict it. This inferiority of the rationalistic level in founding belief is just as manifest when rationalism argues for religion as when it argues against it. (P. 58) https://csrs.nd.edu/assets/59930/williams_1902.pdf

Though from a rational perspective faith may seem like delusion, it is also a form of power, a means to firm up resolve and psychologically brace for the

future. Contrary to rationalism, it is this inner subjective feeling of faith which most powerfully impresses itself on the willingness to believe . Sometimes the willingness to believe is more effective than rationally grappling with uncertainty.

The Nature of Opinion

knowledge versus Opinion: A Matter of Degree or Kind?

"Opinion is like a pendulum and obeys the same law. If it goes past the centre of gravity on one side, it must go a like distance on the other; and it is only after a certain time that it finds the true point at which it can remain at rest."

-Arthur Schopenhauer

Opinions are odd epistemic specimens. They are below knowledge, yet the degree to which they might be considered ignorant varies. Their lack of incontrovertible validity is typically attributed to their subjectivity, their value-driven nature. The analysis of opinion is valuable to the theory of ignorance because of the vagueness of their epistemic status and the fact that they are quite prevalent in everyday life.

Are you entitled to your own opinion?

If I agree with your opinion it does not mean it is correct, only that we share the same opinion. What opinions one might hold are seemingly arbitrary and dependent on selective values, but knowledge is fixed to the one way the world is. knowledge is objective, opinions subjective. The difference is between what has been called direction of fit. Opinions are a world-to-mind direction of fit, which means the world is judged in terms of how it should conform to our priorities. knowledge has a mind-to-world direction of fit: the mind must accomodate the way the world is and determine its factual evaluations.

It might make sense to classify opinions as a subspecies of positive ignorance but this would not be entirely accurate. The view that opinions are arbitrary however, that they are in some sense *beyond reason* is no more firmly grounded than any opinion itself. There ought to be facts about opinions, As James held, there is a *psychology of opinion*, an underlying basis in facts from

which they spring. Opinions are a type of belief, and beliefs are explicable. Somehow subjectivity is considered to be untouchable—just because an attitude or belief is subjective disqualifies it from causal analysis. Opinions are typically thus seen as having a kind of independent existence, such that when an opinion is the only valid form of interpretation of an event, your opinion is as good as mine and if I think you are wrong it is only because your opinion disagrees with mine rather than external evidence. Values are apparently arbitrary, and a value can not be true but can only be. Both knowledge and ignorance come into contract through belief. A consequence of this is that one can persist in maintaining an opinion even knowing that the facts contradict it. Opinions answer to nothing but the values that govern them.

To understand opinion we must at first understand the logical structure of belief. Beliefs are rarely freestanding. To believe one thing is always to believe another and for them to establish a network of intwined mutual reliances. What one values is also dependent on what one believes and knows. The entire system must be analyzed as a global state on which local, singular beliefs are conditioned and vice versa. Any belief or belief system is necessarily bounded; it comes to an end somewhere and has a finite number of belief-nodes. A belief system can be thought of as decomposing into a set of propositions, and propositional attitudes related to each of those assertions about reality or mental content. However, stripped of its relations to other beliefs, a belief cease to exist, it lacks any context. I cannot believe Abraham Lincoln was the 16th President of the United States without first believing that the United States exists, that 15 others came before him, that a person named Abraham Lincoln was born, an so on. Bearing in mind the limit associated with any belief system—we arrive at what might best be visually described as a shell of uncertainty encompassing it. Every belief, question, and epistemic state is necessarily set against a region of uncertainty into which it tapers off into obscurity. The *contained, limited* nature of belief sets against it the world of the undefined and the uncertain.

There is always more of the world outside of the belief than inside the belief. Hence Hamlet's, "there is more in heaven and earth, Horatio, than is dreamed of in your philosophy." What the belief represents is less than all the possible

considerations and points of relevance in the world to which it could apply. Seen this way, the very act of being aware implies a degree of ignorance, because the possession of a mental state or belief at a given time means the exclusion of another. Given that one's mental capacity is occupied by that state at that time, at that time one is ignorant of everything else.

In a strictly formulaic way of putting it: a belief contains less information than parts of the world it attempts to inwardly map, describe, assimilate, accomodate, control, regulate or represent, or any of the things we believe beliefs can accomplish. By necessity, of course any belief is then also less informationally rich than the entirety of the world. Further, the more beliefs are connected in a system, the more fragile their truth status, through the simple fact of probability that the more assumptions one makes the greater the odds that at least one of those assumptions is false. Those behind Hindenburg disaster in which a passenger carrying zeppelin crashed burning into the ground believed correctly that such an apparatus would float, that hydrogen is lighter than air, that it could carry so many people. What they did not believe, or what they never induced in the first place, is that an electric spark might ignite the highly flammable hydrogen keeping the ship afloat. They believed on the contrary that the design of the ship was sound.

While the Bayesian brain as has been addressed in the literature focuses on the most minimal and simplified phenomena such as visual edge detection for reasons of experimental feasibility, a natural extension is that these same principles undergird higher level cerebral functions such as global belief formation. If this transitive property between micro and macro-processing does exist between the two levels of description then the higher level functions can operate independently from strictly sensory, empirical input. In other words, the low level sensory predictions are explicitly empirical. They are entrained by external stimuli. Empirical observation testing of the outer world is a deeply rooted feature of the most basic sensory function of the brain. But on a higher level where beliefs and other complex representations are instantiated, the general Bayesian architecture remains intact despite it drifting away from the strictly empirical character. Beliefs are not always grounded by observation in the

sense that perception is. Beliefs can depend on other beliefs. We have to go out of our way to train ourselves to think scientifically to overcome this subjectivity. Otherwise, that's where opinion sprouts up.

When it comes to a person's beliefs and opinions we consider them to be subjective because they can be freely formed to some degree independently from facts. They only strictly depend on the prior beliefs and the pressure and influence they exert on the formation of new beliefs.

Subjectivity is grounded by facts which provide the context around which beliefs and opinions are shaped. As James emphasizes in *The Will to Believe* opinions are not actually arbitrary and formed on a whim. What we call subjective is not shrouded by unanalyzable magic. Nothing is strictly subjective, as beliefs in themselves are facts among others in the world—though they may not themselves *be* factual. That someone has an opinion is a fact, and there are facts explaining how they came to develop that opinion. Opinions come about *for* reasons though they do not come about *by* reason; people don't value what they do for capriciously: they are structurally constituted by what existentialists call facticity: the sum total of all the social and cultural and situational realities impinging upon them.

Epistemic relativism claims that there is no such thing as absolute knowledge and that what knowledge exists is contextual and conditioned. The same cannot be said about ignorance, because in relation to the fullness of ignorance everyone is ultimately the same in their epistemic finitude. Which is to say, there is a feeling that, within the fullness of space and time, everywhere in the past, present and future, the sum of knowledge may vary but the absolute value of ignorance stays the same. Everyone who is negatively ignorant about a fact is ignorant in precisely the same way, there is no cultural variation to not knowing. knowledge can be added but logically everything one didn't know was already there. Another way to put this is that absolute ignorance is inertial while knowledge "moves" relative to its frame of reference. If epistemology is conceived as how I stress it should, there is no such thing as strict epistemic relativism. knowledge is indeed relative, but not to other knowledge. It is relative only in its variable and fractional relation to total ignorance, as the numerator to

ignorance's denominator.

Opinions are closer to ignorance than they are knowledge, because they are denied the possibly of full truth. An opinion is said to obtain when a belief can be neither true or false but merely good or bad, agreeable or disagreeable with respect to another opinion. There is nothing subjective about the fact of this relativity; opinions are explainable in terms of facts about a given individual's psychology and cultural context. Opinions also have a tendency to evaluate the quality of an object as better or worse. Often, a population's opinions tend to strongly correlate on whether this or that was a good movie or such and such was a great musician. Though there still may be—and assuredly will be—deviants from this conventional wisdom, one wonders what exactly the majority sees in it which they don't. It is this departure, one might argue, which differentiates opinions from what are called social facts. Where everyone is in agreement there are no opinions, only consensus. Opinions, therefore, are inherently conflictive and only arise out of oppositional relationships with out opinions. Opinion emerges from differentiation.

All this begs the question whether objects truly vary in quality or relative utility. Of course there must always be an eye of the beholder to evaluate the object, but the point here is whether there is regular structure across observers that consistently predict statistically normal judgements. Evaluations of something as better or worse presumes difference. Such objects are not what economists call *fungible*, replaceable with another object of its type without any detriment in utility like replacing one mass produced bottle of the same brand of beer for another. When comparing job candidates, the more qualified candidate is typically preferred. Are they the better candidate or is this merely a prejudice? It is not quite an opinion, because the hirer can supply some justification for their decision, but neither can they truly claim to know the more qualified candidate is by necessity better. The unqualified candidate could be a natural savant at the role, who simply hasn't gotten their career started, or the qualified candidate could have gotten all their jobs through nepotistic connections rather than merit. A standard of comparative evaluation presupposes some criteria against which the different comparators are to be judged. It is this selection of the criteria which

is arbitrary in the sense of variable, however *within* that operationally defined context, there are indeed facts as to what is better or worse. Once a standard has been defined, suddenly opinions have the chance to graduate to knowledge. Once someone establishes a law that says "you must be at least 18 years old to vote", the opinion that only legal adults should be allowed to vote is replaced by a social fact. Criteria, even if self-imposed, suddenly elevate opinions to the status of relative knowledge with respect to it.

Museums of fine art do not showcase the immature doodles of children, and neither do literary magazines publish the scribbles of 3rd graders learning to write for the first time. Are these selective preferences not brought about by some force of necessity? Is da Vinci not in fact a better painter than a dabbling amateur, and if not, from where does he derive his acclaim? One can have the view that value judgements are all subjective opinions, but this leaves you inconsistent. How can you know that given the lack of determine facts about that which you speak? Or, one can say it is a fact that variable examples of a type differ in quality, that da Vinci truly is a better painter than dabbler without control over the medium, technique or inspiration. And you can either be attuned to this fact or not. If you think scribbles are better than the Mona Lisa you are factually incorrect or both incorrect and biased. One who favors the dilettante over the master is probably biased by some motivated reasoning, it could be their mother perhaps, or they may have quirky tastes. Suppose then there is a world where everyone has the same opinion for a given case. Does such an opinion cease being one or is it instead a fact of that consensus reality?

The sentiment remains that some opinions are truer than others, that there can be variably informed or ignorant opinions, especially given evaluations of quality according to well specified criteria. Even allowing that there are different tastes between people, there is a finite range for such tastes, and there are sufficient reasons for why one individual might prefer it this way rather than that. Why else would we trust experts in creative fields to instruct and grade art students if it were all matter of taste? Their judgements are not as absolute as those of the engineer's, but they are able to specify reasons for their gradations of relative quality. And yet opinions are inferior to knowledge or even just

unprovable but well reasoned arguments. They default to ignorance in a special way because they are about what isn't or cannot be known, not in the sense that there is a fact that is out of reach but because there is no fact. If there was a fact to be known the formation of an opinion would be preempted by the knowledge of it. Any situation of fact in which either an opinion *or* knowledge could hold, it is not that all opinions are equal it's that all are invalid. In this situation there is no reason to choose opinion over knowledge except save for emotional stubbornness. It is only in conditions where knowledge is either impossible, impracticable, or too effortful to achieve, that opinions have any standing.

Some opinions, however, seem more like knowledge that has yet to grow up. These are more like conjecture. Add some measurements, impose evaluation criteria, refine the arguments, find evidence, accumulate data and suddenly an opinion inches toward knowledge and away from ignorance. Were this truly the case, then this destabilizes the very basis for opining. Could it be that opinions attempt to fill the space of the unknown, unknowns at least in theory it might be possible to know but which are out of reach? Or do opinions attempt to compensate for the unknowable by filling in what cannot be known with at least some suppositional mental content? It may indeed be a measurable, definite, provable fact that one socioeconomic system is "better" than the other, or that one policy causes more harm than good, given some precise operational definition of what it means to be a better or worse social system. It feels epistemologically threatening that such important topics cannot be proven with finality. The criteria of evaluation selected to interpret what is true and false, do not themselves have criteria and therefore any knowledge that might be gleaned from them must be merely relative to the limitations of the theory's assumptions. Reason chafes against such ambiguity, because it feels that it should be able to "write the book of the world" in Sider's phrasing. It should be possible to match the correct set of concepts with reality ([https://www.amazon.com/Writing-Book-World-Theodore-Sider-ebook/dp/B00NLB62WU]). Mill's utilitarianism is an attempt to quantify the objective value of policies in terms of whether they bring the maximum happiness to the majority. Even then holes of vagueness riddle this doctrine. How to measure happiness? How to be sure it has reached a

maximum? Why the majority? How to poll the majority in a fully descriptive way to see if they actually benefitted? Who decides what will benefit the majority if not the majority? How to get them to agree? Such utilitarian thinking is laden with assumptions, layering ambiguity on top of ambiguity. Beyond the questions just listed, there is the additional problem of whether society is best thought of as a machine, consisting of variables to be measured, maxima and minima, and parameters to be optimized, and adjusted accordingly through the application of neutral and dispassionate technocratic efficiency. This is opposed to viewing a society as consisting of flesh and blood people, each expressing a degree of individuality and autonomy but bound by a generic human nature and social ties, and chafing against or cooperating with one another in in variable states of contentment. Happiness is not a simple quantity to be weighed or counted. Once again, in the effort to understand the complex is simplified to the point that what is understood is not the actual object but a toy model of it that is by no means built to scale.

The reason for the infeasibility of political knowledge however is not that it is truly indemonstrable in principle, but that it is beyond all practicality to formulate valid models with all the parameters and complexity of real world economic or political systems and run them against each other. History only happens once, and social experiments are particularly odious because of the ethical costs of using people as lab rats. Social experiments such as fascism and communism that utterly upended their former societies differed from what came before them in that they projected a supposedly complete, comprehensive ideology of a supposedly ideal social structure over the flawed, inertial "resting state" of prior historical continuities. Ancien régimes reflect a sort of unconscious, accretive Jamesian pattern of habits, the traditionalistic "way of the world" as it had built up generationally, rather than an explicit ideological project which consciously and deliberately seeks to map a set of values to the world. Inversely, revolutionary movements often fail or achieve counterproductive results their every embrace of novelty and possibility by necessity thrusts them into the trajectory of the unknown. The French Revolution, which sought to abolish all kings, and promote liberty and equality,

led to the brutal repression of the Terror and ultimately to the dictatorship of Napoleon. The status quo may be bad, but its continued existence is at least proof of its stability.

Because a subjective core of feeling and value is at the heart of all these questions, deciding once and for all whether abortion for example is acceptable, or that democracy is inherently better than monarchy, or that the rich should be taxed more, in the same way deciding that 2+2=4 is true seems untenable. The main reason for this is intersubjectivity. From the standpoint of the monarch about to lose his birthright, or the rich about to be taxed, idealistic reflections on whether it is socially optimal to abolish monarchism or progressively tax the rich fall by the wayside. Because these people are part of the very system which is under analysis, how they react to these policy suggestions in the end contribute to how they play out. Yet what we have here is arguably not world ruled by arbitrary opinions, but an overly complex, inconsistent and conflictive social organization in which some parts refute the aims of others rather than acting in organic unison.

Moral sentiments somehow mysteriously escape logical judgements, but it feels weird to state that it cannot be known what is right or wrong and that we are ignorant of moral truths, given the intense potency of conviction with regards to the matter. This rejection of the knowledge of moral truths would also commit us to moral relativism or even nihilism. The sadist would be just as justified as the saint because neither has any ethical basis for their actions. When someone holds what are widely seen as morally reprehensible views, those with a conscience recoil against the insinuation that it ultimately comes down to "their words against mine." At the same time, a morally corrupted society can view justice as wrong-headed. Absurdly, scientific racists in the antebellum American South proposed a fake clinical diagnosis, drapetomania, to explain the supposedly mentally ill behavior of slaves who sought to escape captivity. Not only does variation in values and perspective color opinion, but so does power.

Those with a strong sense of justice typically feel in it as overwhelming force that transcends whatever small-minded attitudes of those who embrace morally unjust positions, which always aligns with moral ignorance. This moral force has

the insistency of instinct, it seems to work itself up in us with an involuntary emotional compulsion. In the words of preacher and abolitionist Theodore Parker that were later paraphrased and immortalized by Martin Luther King,

> We cannot understand the moral Universe. The arc is a long one, and our eyes reach but a little way; we cannot calculate the curve and complete the figure by the experience of sight; but we can divine it by conscience, and we surely know that it bends toward justice. Justice will not fail, though wickedness appears strong, and has on its side the armies and thrones of power, the riches and the glory of the world, and though poor men crouch down in despair. Justice will not fail and perish out from the world of men, nor will what is really wrong and contrary to God's real law of justice continually endure.

This passage is remarkable for its conceptual blend of logical and moral comparisons. Reason, metaphorically rendered as geometrical figures, cannot measure justice. Morality cannot be interpreted as a rational geometry, Parker says, it cannot be truly known empirically. It is not reducible to a predictable formula. Yet justice nevertheless has an inexorable reality, a visceral truth, which operates according to God's and by extension Nature's law. Though it may be ineffable, there really is a "reason beyond reason" to justice. Parker suggests that there is an immanent, trans-empirical logic to the conscience, which is governed not the laws of man but the rule of God. This moral truth cannot be articulated by intellectual reasoning, but it can be felt in the conscience, which, almost as Adam Smith would have it, perceives moral sentiments almost as if by a sixth sense. That feeling is not just a fickle mood but evidence of the conscience transcending the discord of moral ignorance and coming into moral wisdom.

Why certain topics of discussion seem to be designated as opinion-only is not therefore strictly because knowledge is impossible in those domains. It is because overwhelming intersubjective complexity creates unmanageably elaborate chains of relations between people and their beliefs, people who often have conflicting and competing interests and priorities. Opinions cannot be entirely arbitrary, because they represent the cognitive evaluation of an object by an agent with a

definite nature in relation to that object. While the fact that opinions are acts of an agent's evaluation alone is not sufficient to promote the opinion to the status of a fact about the object, it demonstrates that the object was capable of inducing such a mental state in at least one subject with such a constitution. Opinions are caused by facts though they may not be factual. As such opinions are fuzzy epistemic constructs that straddle the boundary between ignorance and knowledge, never quite settling into either category, but able to lean more towards the one or the other. To expand on the matter of the nature of opinion, the colloquially recognized notion of willful ignorance, must also be better understood.

Willful Ignorance as the Will to Believe

No one needs reminding that people can be irrational. Irrational does not always mean bad. For example, one can argue that falling "madly" in love is positive experience, at least up to a limit. Or one can feel an obsessive passion for a pursuit which drives one to excel at it. Irrationality in its essence might even be taken for what truly makes one fully human. Perfect rationality is a kind of confinement. In the Nicomachean Ethics Aristotle distinguishes between the rational and irrational soul. The rational soul is occupied by the faculties of reason and is associated with the mental intellect, whereas the irrational soul is bound up with appetitive bodily desires. Yet when we think of a being with too much of a "rational soul" we end up with something soulless: the mechanistic logic of the insect or the insensible calculative dispassion of a computer. Typically it is supposed that it is humanity's capacity for logic and reason which sets it apart from the brutes, but if anything animal nature is ruthlessly rational and little in the world of nonhuman biology is arbitrary or whimsical. It is instead the highest of free feeling, drama and tragedy, that rises above the merely fitness-optimized dictates of survival and reproduction and which set humans apart from the rest of the tree of life.

From an epistemological standpoint perfect rationality would mean always preferring truth to falsehood, and never choosing ignorance over knowledge if

the option were clearly presented. Ignorance is often something that harms, an affliction or shortcoming, but it can also at times be a deliberate choice. To embrace truth is in a way to submit to a reality not of one's choosing; if the truth does not align with one's aims, then assent to it comes at a cost of decreased autonomy. So people do not always choose truth over falsehood. *Homo insapiens* intimately embraces ignorance if it is somehow useful, if it allows them to behave in ways that allow them to construct their own rules. People, argued Freud, repress traumas. They entertain comforting illusions and fabricate exaggerated self-images to impress and beguile others. They indulge in narcissistic fantasies, gladly promote not what is true but what benefits them, and actively seek to hide the truth from themselves if it risks popping the bubble of epistemic security to which they are so highly emotionally invested.

The phrase "willful ignorance" denotes this type of self-defensive, or self-serving, opportunistic denial of knowledge. Yet more than a denial of knowledge, it can be the drive for positive ignorance, the desire to construct a world of wishful thinking which is not coerced by the current facts as they stand. Willful ignorance is in the last a consequence a subtype of what James called the *will to believe*. James centers volitional agency at the core of belief, rather than pushing the excessively rationalistic doctrine that mandates that rationally justified belief must always conform to rigid objective facts. James is not concerned for example with situations in which one believes one is well when sick, or that two dollars in one's pocket are actually one hundred. Such decision of such a truth according to James is what he calls a *forced option*, which is to say, an option for which there is only one unavoidable choice. There are however situations in which the determination of the facts is either vague or undecided, or which the outcome is potentially influenceable by an agent's actions and beliefs. In such cases James favors trans-rational, as opposed to irrational volitional desire as the decisive factor in the selection of one option over another. Human beings, in other words, *make* facts, they are not just *made* *by* facts. Such desired outcomes he calls *live* as opposed to *dead* hypotheses. Such options elicit a willingness to act, and thus one is motivated to believe in them. Viewed pragmatically, then, opinions are driven by the logic of *motivation* rather than

the logic of truth. In this way James comprehends belief as a means to determine what one wants even in the face of ignorance as to the facts. It is a view that appreciates a participatory universe conception, in which human agency actively shapes potential truths that eventually become true rather than passively monitoring the world in the expectation that its facts are part of some pre-established harmony.

At first it seems implausible how people can operationalize, or at times even weaponize ignorance. Ignorance it is commonly assumed is always a deficit. knowledge, after all, is power. However, as James states, this does not deter one from belief. "As a rule we disbelieve all facts and theories for which we have no use." In this pursuit of ends, "pure insight and logic, whatever they might do ideally, are not the only things that really do produce our creeds." This brings James to the crux of his analysis:

Our passional nature must, and lawfully may, decide an option between propositions, whenever it its a genuine option that *cannot by its nature be decided on intellectual grounds*; for to say, under such circumstances, 'Do not decide but leave the question open,' is itself a passional decision, just like deciding 'yes' or 'no' and is attended with the same risk of losing the truth.

In summary, assent to a proposition and the attendant commitment to a course of action *without rational grounds * is permitted if the dilemma posed is *undecidable from the point of view of all available information*\[emphasis added]

The emphasized clauses in this passage capture James's implicit grasp of the structuring role of ignorance in decision making and the epistemic evaluation of vagueness. For James even some of the most basic epistemic propositions, such as "there is truth and the mind can obtain it" fall under such an act of passional will. Indeed it is not clear even after thousands of years of debate what exactly truth is. The truth about truth assumes itself; we'd have to know what truth is already to fix an answer to the question we have about it. If we waited for a perfectly clarified and accurate idea before proceeding nothing would get done. From the "point of view of available information" there is no ultimate, ultimate signal that what is available is sufficient or complete—until after the fact. There is

a class of questions that stand in perpetual want of of a perfect and complete answer. “The absolutists… say we not only can attain to knowing truth, but we can *know when* we have attained knowing it.” This is not necessarily the case, save perhaps for certain trivial toy examples or contrived deductive logics, because, “to know one thing, and to know for certain *that* we know it is another.”

James has a grasp of positive ignorance and the mind’s proclivity towards it. “Biologically considered, our minds are as ready to grind out falsehood as veracity.” And so it hinges on “volitional nature” to tip the scales in the absence of perfect information, to proceed with a *working* hypothesis that allows the business to proceed. As the temporality of information constraints shows, time pressure is one of the primary sources of structural ignorance. James is an empiricist, and so holds dispassionate scientific impartiality up as the ideal of rational judgement, but concedes that not every situation allows for the controlled conditions which make scientific reasoning feasible. Though a surgeon is trained to think with empirical discipline, conducting a careful, risky and time sensitive operation affords opportunity to experiment or replicate tests. Moral questions in particular are a class where evidential proof is not forthcoming. Indeed the question of moral truths is ambiguous. It is not clear if in the real world, ethical quandaries are treated as “problems” to be solved, as opposed to as sort of blunt, color-toned emotional states devoid of calculus. Moral uncertainty is not the same as the regular perceptual and cognitive kind associated with the apprehension of the sensible world. It is not based in undecided information entropy that establishes an imbalanced gradient between an agent and the environment which is successively diminished through proactive probing behaviors and sensorimotor investigation. Moral truths and the ambiguity about them are not plainly informational in character.

When the rational intellect stops to contemplate inanimate nature, the object of thought presents itself as blunt. That is, inanimate nature lacks intentionality and moral standing. The debate around moral truths, and consequently knowledge or ignorance about them, goes back to the so called fact-value distinction. In this case there is a confusion between (factually) true and

(morally) right, or false and wrong. Is it a fact that it is wrong to punish an innocent person, or that it is right to punish the guilty? There certainly is a strong consensus regarding these issues. As DeNicola writes, the connection between factuality and moral judgement depends on one's relationship to the fact:

If the content of a belief is judged morally wrong, it is also thought to be false. The belief that one race is less than fully human is not only a morally repugnant, racist tenet; it is also thought to be a false claim – though not by the believer. The falsity of a belief is a necessary but not sufficient condition for a belief to be morally wrong; neither is the ugliness of the content sufficient for a belief to be morally wrong. Alas, there are indeed morally repugnant truths, but it is not the believing that makes them so. Their moral ugliness is embedded in the world, not in one's belief about the world.

(https://aeon.co/ideas/you-dont-have-a-right-to-believe-whatever-you-want-to)

According to pragmatism, the feeling of rightness or wrongness here is a fact in itself, an inner necessity that makes some options live and momentous and others dead and trivial. In James parlance, the option to punish the innocent is not forced in the same way that wanting to survive and get out of the way of a speeding train is, but it does serve as an affront to the moral passions. This brings us to the larger, main point.

Typically it is assumed that a valid belief must have what's called a world-to-mind direction of fit. Meaning that the way the world is ought to condition the truth of the belief. On the contrary, a desire is said to have a mind-to-world direction of fit. It is expected that when one desires a result that the world ought to conform to the desire and not the other way around. Yet these distinctions between belief and desire break down once the categories are mixed. What is the proper direction of fit when one desires to believe?

For one who chooses what might be considered an ignorant position for their own benefit , they are doing so because it is a live option for them. There's a perceived payoff. Beliefs are facts, though they may or may not be factual. They are a thing—granted an abstraction—that can influence the world and as such reshape it to resemble the world the belief projects . Such motivated reasoning

does not submit to the world of facts, but seeks to warp the world of facts until it conforms more to its designs. In this scenario, the positively and willfully ignorant is actually pragmatically knowledgeable, they just prioritize their own view of things as a fact over the alternatives and make it so. Someone who is willfully ignorant, that is, who volunteers to belive something that does not fit the way the world is at the time they form the belief, may nevertheless desire the world the belief projects and thus change the facts as they stand. Ignorance therefore can be in some cases useful, because a believer who stops at the apparent facts and does not dare to contest them surrenders their agency to external forces.

For these reasons willful ignorance is really the will to believe, a wish to make a value a fact. It is surely the case that some values have won the struggle to become facts. In the final analysis, willful ignorance makes the most sense if conceived from a pragmatist angle. Beliefs are tools and at times it can be advantageous to believe wrongly.

Cultural, Artifactual and Generational Ignorance

The notion of ignorance production lends itself naturally to the idea that ignorance, especially of the positive variety, can be constructed as an artifact that leaves its mark on the world as an effectual cause . Just as the consequences of accurate knowledge affect the world by ordering it more rationally, so does a kind of entropic mental kipple muddy things up. Like garbage in the streets or the near omnipresence of pollution in the modern age, wrong ideas and wrong actions spill aimlessly into both mental and real space, scrambling it and making it a more irrational place. What is pernicious about these patterns of ignorance is that by virtue of them being undetected products of unknowing or misbelief, they often pass by unnoticed until the damage has been done. Climate denialism and the unforeseen consequences of fossil fuels is an ominous example of such a tendency. The so called law of unintended consequences, goes to show that the shelf-life is often shorter than expected and what made perfect sense decades ago and continues to this day can be foolish in the present even though it has over

time become stealthily normalized for modern sensibilities. At one time the avid use of plastics seemed like industrial miracle in materials production. Now it pervades our oceans, piles up in devastatingly vast landfills and gathers worrisomely in our lymph nodes.

The question of cultural ignorance can be approached from several angles. Just as humanity inherits the insights and good decisions of their ancestors, so too their errors. Just as the past flows in general into the present, those aspects of the past that consisted of the sum of human action channel into the present in a more concentrated and communicable form as culture. Legacies are double edged. knowledge and art are left over, no doubt, but so is the ignorance and the messes of previous generations. Not only are the cultural prejudices of previous generations bequeathed to their descendants, but so is the aftereffects of their mistakes and poor decisions, or even just the consequences of their limited knowledge applied to a world that has since changed and so lingers on unsuited to the conditions of today.

The current civilizational moment is path dependent on the sum of decisions made by every person who ever lived. However, given that both knowledge and ignorance is transmissible, it behooves present generations to act as a kind of sieve to detect and filter the mistakes of the past. New generations collect what is transmitted from the previous, all their solutions and knowledge, but also all the unfixed problems and even those they caused, and all the errors and backwardness they left unquestioned. This legacy is immutable and forms a backdrop of culturally transmitted knowledge and ignorance. It is what present generations must "work with", history as it happened is their starting point. These generational transmissions become the milieu which the descendants either recapitulate or revolt against. This puts new generations in a tight spot. They must inevitably draw some context from their predecessors, and at first as children this is obligatory as their only chance at survival is to listen to their elders. But it is also in their nature to transgress this cultural past, overcome it, revolutionize or evolutionize beyond it. If anything this transgressiveness is expected and desired of them—to a limit. The older generations stake their hopes in that their offspring will improve what they could not and contribute what they

didn't. By the same token, "business as usual" can only continue insofar as the future resembles the past. Ultimately, cultures change because the world changes.

Broadly construed, a few perspectives can be offered about the epistemic value of cultural inheritance:

1. It represents the well meaning efforts and good intentions of the previous generation of whom the next are their descendants for whom they are concerned . It is the result of all that they grappled with, accomplished, created and built as they struggled as imperfect beings through the joys and pains of the human condition. This legacy is value-laden and attempts to communicate messages of lived wisdom and instruction to posterity. Though flawed, it speaks to the human spirit for good or ill, in all its contradictions, strengths and frailties. The past generation strives to impart more of its knowledge than its willful ignorance.

2. It signifies merely tradition for tradition's sake, a rigid concrescence of habits that perpetuate as it were autonomically. This chain of habits transmits by imitation and memetic mirroring. There are no intentions behind the generational inheritance, what transmits is for the most part incidental. It's whatever sticks. You get what you get. Those in the past were concerned with living their lives in the moment no more and no more less than you are now. When they passed they left it how it was. Previous generations impart more or less equal parts knowledge and ignorance.

3. The past is meant to be transvalued and overcome. The unresolved work of the previous generations is merely a stepping stone or worse, a hazard to be disassembled and cleared away. It is deeply erroneous, a problem to be solved, and it is the next generation's responsibility to fix up the mess the previous, in their negligence and shortsightedness, selfishly made. Earlier generations mostly leave behind their failures, their ignorance. After all, generally speaking the more time goes on, the more becomes known, so newer generations are privileged to rewrite the rules to accommodate current conditions.

These three perspectives offer different judgements about cultural legacies

ranging from the most charitable and optimistic, to neutral and almost mechanistic, to the cynical and pessimistic. The first case suggests that the prior generation is selective about what they pass on, seeking to leave behind what has value, cultivating what they leave behind. Whatever ignorance they impart is what they failed to address or discover, or which they could not overcome. Humanity seeks to build itself up intergenerationally and this is perhaps an extension of the parental instinct. We stand on the shoulders of giants. The second case is an expression of brute memetic virality. Only that which propagates does for no reason other than it can. There is little to be done to control the power of cultural inertia, and a collection of accumulated habits passes through each generation like a force of nature without any conscious effort able to intervene and stop it.

The third option conveys that past generations for the most part leave things more polluted, used up, and confused than they found it. We inherit their mistakes. Those who came before us serve only as an example of what not to do. Progress is when each new generation sheds the ignorance of the previous.

All things considered, the truth is probably some mix of these depictions. No one however can deny the ponderous momentum of cultural habit referenced by the second. In chapter 8, The Laws of Habit, of William James's monumental *Principles of Psychology*, he describes habit as an inexorable, dominating force which nearly enslaves the whole of humanity and keeps civilization wound like a clock. "All our life, so far as it has definite form, is but a mass of habits,—practical, emotional, and intellectual,—systematically organized for our weal or woe, and bearing us irresistibly toward our destiny, whatever the latter may be." It may be the case that rather than deviating too far from the known, an innate behavioral conservatism tends to take precedence in older generations, as they cling to what works and remember the sting received for straying too far out of line. This grim deterministic view of habit can be extended to suggest that the parent essentially aspires to parasitize the child with their same structure of habits, to whatever degree of success. The younger generation however often seeks to test the boundaries, as is the nature of youth. Each generation after a fashion must repeat the same mistakes for the pattern to reestablish itself anew.

Surely each generation passes along more than just its knowledge to the next. Alluding to the cultural and generational encoding of such ignorance, Schopenhauer darkly observed, "While illusion distorts reality for a moment, error can reign for a millennia in abstractions, throw its iron yoke over whole peoples and stifle the noblest impulses of humanity; those it cannot deceive are left in chains by those it has, by its slaves." The Romans, like many ancient societies, traced their origins to mythical figures, namely Aeneas, prince of Troy. Many young children simply believe what they are told as if by reflex. It has even been speculated that children's seemingly innate uncritical receptivity to information is an evolutionary trait. A prehistoric child who doubts what their parents tell them about the snake is probably less likely to see adulthood than one who faithfully assumes the truth of what they are told. Skeptics don't survive long in the wild. Children tend to be gullible because this may maximize the rate of learning. (Dawkins R. Viruses of the mind. In: Dahlbom B, editor. Dennett and his critics: Demystifying mind. Oxford, England: Blackwell; 1993. pp. 13–27.). This is another example of ignorance serving a constructive purpose, albeit one with glaring drawbacks. Parents who celebrate Christmas relish tricking their child into believing in Santa, and generation after generation of children fall for it unfailingly, with rarely a dissenter to be found among their ranks. Parents playful construct such fanciful ignorance in their children because they remembered the magic of believing in Santa themselves as a child, and find such innocence and misplaced reverence cute. Or perhaps they believe in the merits of behavior modification by telling them that Santa will leave coals instead of presents if they misbehave. Regardless of its harmlessness, such a myth is a case of generationally transmitted deception, agniogenesis. And the same goes for all the mythologies and invented heroic pasts which elders lecture to the young, instilling in them a grander and more nobler vision of their history than what really was in the hopes that belief in this ideal will drive the young to excel and dream.

Such cultural fictions Freud wrote about in his book *The Future of an Illusion*. These are, "certain dogmas, assertions about facts and conditions of external and internal reality which tells one something that one has not oneself

discovered, and which claim one should give them credence." He later continues with a striking observation: "Thus we arrive at the singular conclusion that of all the information passed by our cultural assets it is precisely the elements which might be of greatest importance to us and which have the task of solving the riddles of the universe and reconciling us to the suffering of life— it is precisely those elements that are the least well authenticated of any." It is the objects of faith, the ideals of the spirit, which seem to have the most durable power. It is not with any prejudice that I call faith an object of ignorance. It's not that to be faithful is to be ignorant, but faith is a response to ineluctable ignorance. Faith picks up where knowledge falls. Freud means to say that what is most communicated by cultural artifacts is that which falls under the least scrutiny, those data which, if they were doubted too harshly by the young, would destroy the psychic and social integrity of the culture.

What would become of a civilization where children were taught only their parent's doubts? Parents naturally assume an authoritative epistemic posture when instructing their children even though they themselves may not know with certainty in the depths of their minds. They must provide a model of epistemic security which serves as a bulwark for their child's vulnerable and systematic ignorance. What would come of it if they simply told their children, "I don't know. Figure it out yourself"? The children can't. Generational ignorance is therefore smuggled in through generational knowledge transmission. The young must be taught what the old never finished figuring out.

Next to the category of generational ignorance we might add artifactual ignorance. Just as knowledge may be encoded in external objects such as diagrams, software, and books, so too can ignorance. The wrong ideas are transmissible not only through words and the notions they convey but also the material culture, the practices and constructions which mediate life. To the degree that people are clueless that the artifacts they use are not the best way to do something, that object is enclosed within a boundary of ignorance. Under this definition nearly everything is an artifact of ignorance, however, because nearly everything which has a used can be improved. Yet some objects are more overtly than others.

In the extended cognition literature external apparatus such as tools, technologies, written language and so on are thought about as enhancing baseline, brain-centric processing. We assemble environments that bolster our thinking and support memory: laboratories, libraries, universities, databases. What about the opposite? How do our constructed environments encode our mistakes and follies? From one vantage point it would seem to contradict the pragmatic criterion of truth that a constructed object could embody ignorance. Just as how the construction of a bridge demonstrates a proof of the assumptions that went into it, so too anything that lasts. A bridge constructed in ignorance is a bridge that collapses. This is not out of step with what I'm getting at. Some knowledge is perishable, it contains a time bomb of ignorance that slowly ticks down. Such knowledge has an exportation date; what is true then is not necessarily true now or later if the fact is contingent. Such is the inevitable when knowledge is fitted to a situational context that is temporary. Many artifacts are constructed to work only within the horizon of possibilities which their creators could envision in their time. After they exceed these horizons they become not a solution but a problem. Fossil fuel driven industry seemed like a great idea at the time, and while it still gets everything done, there may very well come a time when it's not so great an idea. Other artifacts are constructed out of ignorance whole cloth.

Medical treatments of the past strike us now as absurd, tinctures infused with cocaine, the quicksilver potions for eternal life imbibed by ancient Chinese emperors, the bloodletting and frontal lobe lobotomies that did more harm than good. Yet there were clear rationales, however mistaken, to create these false practices and items. They were invented out of the pursuit of knowledge. There was ignorant thinking behind them, but ignorant thinking produces ignorant objects. In a somewhat related sense, the material culture can display ignorance of choice rather than artifactual output. According to archeological evidence the Mayans of Central America had no practical concept of the wheel for transport, but did have what are thought to be wheeled children's toys. ([https://www.scientificamerican.com/article/why-it-took-so-long-to-inv/]) Perhaps they were satisfied with their existing means of transportation, or perhaps the

practical value of wheeled carts and similar devices never occurred to them. Arguably given the nature of the earth's gravity and the laws of geometry, wheels are an optimal means of transportation on flat surfaces that cultures have therefore tended to independently discover. The ancient Romans used a clunky and impractical numeral system which makes doing advanced mathematics virtually impossible. It even lacked a symbol for zero. The ancient Greeks used a letter-based mathematical notation, which while far less flexible than the Hindu-Arabic notation used ubiquitously today, had greater expressive power than the Roman. The Romans, having conquered Greece, undoubtedly knew of this Greek system, but out of cultural hubris must have preferred their own idiosyncratic symbolism. A Roman soldier murdered the brilliant Archimedes for his disobedience. According to the story, Archimedes refused to comply with the demands of the soldier who barged into his study and instead stubbornly persisted in his work. This one act of a legionary arrogance wasted what would have been a priceless asset for the military engineering of the Roman war machine, denying them the construction of many intelligent artifacts. If you go looking for them, many such cases can be found throughout history.

Only out of scientistic pride can we tell ourselves that modern culture is beyond such shortcomings. There are sure to be numerous similar lacunae or inefficiencies within our civilization's operating system that we simply don't recognize out of habit or are deceived by the illusion that current technological setup is immutably normal. The world of today cannot conceive of a world without cars. Is there a better way to treat cancer other than subjecting the human body to harmful radiation or toxic chemicals? The belief that we have the one and only treatment prevents the discovery of a potentially better solution. (It must be granted that this is not the best example, given that in all likelihood medical professionals recognize that chemotherapy is a satisficing solution to cancer treatment in the absence of any knowledge of a better one.) One can imagine nanotechnology that targets cancerous tissue and eliminates it through some precision means. And so on for all technological norms and solutional practices we take for granted as *the* practices.

Some may at some point in their lives rebel against current understandings. In

doing so they might raise to awareness the possibility that the way things are could be different, but only occasionally do such alternatives manifest. A revolution without the plan to implement an alternative is merely a rebellion, futile thrashing against authority. Is capitalism the truly ultimate template for economics? It prevails because a better system is inconceivable or impractical (or because of strategic ignorance in the elites who profit most from it). The sheer momentum of custom enforces a kind of cultural ignorance, an impenetrable sensation of inflexible normativity. Consensus reality has the power to determine what's socially possible. Eventually many accept that what happens is "the way of the world" or "just how it is." Occasionally an innovation will arise and disrupt the way things are done, which pushes the culture along and reveals the baselessness of the formerly presiding Eternal Truths. But generally the force of habit is not to be underestimated. The risk of straying too far from the familiar becomes a constraint to which discovery and invention is limited. In this way even finite knowledge can be a source of ignorance once it is diminutively truncated to fit ideological strictures.

The manufacture of artifacts of ignorance can be caused by ignorance, can be suboptimal for their purpose as a result of ignorance or poor choice, or they can be made *to cause* ignorance. As detailed in the previous section about ignorance production, objects can be designed to deceive. Halls of mirrors, metaphorical or otherwise, boobytraps, false advertisements, decoys and cooked books—it's a true challenge to exhaust the list of such items. If you don't know better verisimilitude can be mistaken for authenticity. The object is made with the knowledge that some will think it is what it's not. The most effective lies, as we will soon see, are those that trim their exaggerations and distortions to fit the contours of publicly acceptable plausibility. Such insincere false negatives infiltrate our honest conceptions, like moles behind enemy lines that compromise their strategic intelligence.

We are now faced with a technological age with the potential for unbounded artifactual ignorance. Machine learning systems can produce deepfake simulations of startling verisimilitude across media, including audiovisual and textual. Many lack a sort of critical immunity to catch such deceptive artifacts in

action. We are used to thinking that portraits, and especially videos of people are real, and that what is written or drawn has been drafted by human hands. And while photoshopping has gradually increased our skepticism about the veracity of photos, in the past it was much harder to construct video recordings that were uncannily realistic yet fake. The same can be said of natural language processing programs such as GPT-3, which can generate believable texts, usually detectable for a fake only because of an occasional lack of semantic coherence or discrepancy between it and our background knowledge. Otherwise these computer generated texts, with exceptions, syntactically pass the Turing test. This recurring theme of an ignorant thing done intelligently is once more applicable. It is hard to conceive how these remarkably advanced technologies that were the products of ingenious innovation can be used for any other purpose than to deceive and spread ignorance. Such misinformation technology beckons an entirely new level of ignorance manufacturing.

Both generationally and artifactually, ignorance can become entwined with cultural life. Because what we are most deeply ignorant of is unrecognizable, these strands become embedded into the whole fabric of the cultural operating system. Common sense becomes common nonsense. Because people invariably think and understand by implicitly relying on an extended cultural knowledge base, their ignorance is correspondingly abetted by a kind of cultural "ignorance base." New generations inherit this flawed framework. It is also wrong to think that to be ignorant prevents anyone from building anything, that ignorance always leads to a bridge that collapses. Ignorance can lead to suboptimal constructions, and in combination with generational ignorance, can persist such imperfect patterns until it finally occurs to someone that the status quo leaves much to be desired. Cultural ignorance can also exhibit a regulative maintenance role, it can keep things from changing.

Part VII: The Theory of Inquiry

Who really knows?
Who will here proclaim it?
Whence was it produced? Whence is this creation?
The gods came afterwards, with the creation of this universe.
Who then knows whence it has arisen?

— Nasadiya Sukta, concerning the origin of the universe, Rig Veda

As has by this point belabored to exhaustion, knowledge and ignorance cannot be understood without respect to each other. What's missing is an understanding of the interface by which ignorance transforms into knowledge. Lacking this middle term there is no way to relate ignorance and knowledge in a meaningful way. Epistemic agents intentionally behave in such ways as to resolve and diminish their ignorance. They inquire. Without doing so, it is impossible for them to come to know learn outside of mere accidental discovery. Because ignorance comes first, if they cannot come to know anything they cannot know anything. Learning precedes all knowledge and ignorance precedes all learning. Therefore inquiry is the critical link between knowledge and ignorance. Rather than to passively await the accidental resolution of their ignorance, a helplessly leaving it to chance, epistemic agents inquire, investigate, question, search, explore, probe, dissect, and experiment. The awareness of ignorance becomes the motive force for the pursuit of knowledge, without which the only source of knowledge is haphazard trial and error, to have the passive reception of pain and folly be one's sole teacher. Optimal learning does not happen incidentally; it is targeted and intentional, involving the intentional awareness and successive diminishment of ignorance which recedes in inverse proportion to which the progress through the relevant information exhausts those resources.

To tie the relation between knowledge and ignorance together a theory of inquiry must be developed. It is an inevitable requirement if a more complete

epistemology is to be understood. It gives movement and dynamism to knowledge, in the same way that biology is made sensible by evolution which has taught us not to see lifeforms as static types but as entities that continually change with time and have extended, hereditary lineages that determine how they are. Inquiry is the middle ground between knowledge and ignorance, it is the connective tissue that binds them and without which knowledge cannot exist. It draws the circle to a close.

The theory of investigation can be divided into two main categories: *what* investigation is and *how* the unknown is investigated. To understand how something happens, it is necessary at first to know what that something is. The tricky part is that the best way to find out what something is, is to observe how it behaves. This is even more the case when what is under the microscope is a kind of behavior. That said I would argue that investigation is not strictly reducible to investigatory or exploratory behavior, that this behavior is instead an outward expression of an underlying internal logic the skein of which can be untangled from the sensorimotor enaction. For this reason, I will open with what inquiry is before moving on to an overview of some of its mechanisms and methods.

In recent decades a veritable philosophy of questions has taken shape, giving it the vital attention it deserves. In particular, Hintikka's work on what he calls *Socratic Epistemology* is a significant landmark in this new theory of inquiry. However, Hintikka's epistemology, I argue, proceeds without an adequate understanding of how ignorance factors into it. As such it is this critical missing element which as will soon be shown restricts the reach of his analysis. His approach however seems to skirt around certain implicit assumptions which to my mind are at the core of the matter. We will examine Hintikka's work in due course, however it is these implicit ideas that are the next concern. I call this fundamental intrinsic mode of investigation the *inquisitive stance.*

The Inquisitive Stance: Ignorance as Intentional Object

To understand inquiry, the inquirer must first be understood. Inquiry is an act

or at minimum a cognitive relation associated with agency. Inquiry unfolds with respect to a mind and its suite of functionalities and features. characterize the pattern the cognitive system adopts when in an inquiring mood. Something special happens in the mind when we append a question mark to a sentence. This state of inquiry is not explainable in terms of an epistemology lacking in a proper treatment of ignorance for the obvious reason that the cause of the inquiry is not an object of knowledge but of ignorance. I call this mental process the *inquisitive stance* and identify it by several necessary attributes. In short, its defining features include *self-recognition of information inadequacy*, *ignorance as an intentional object*, *epistemic emotions (such as wonder, curiosity)*, *information seeking* and *information equilibrium*.

Taken together, these features comprise an engine of inquiry that functions to zero out a portion of ignorance that has been identified to a such a degree of salience that a question can be formed about it and appropriate information sources sought for it. This adoption of a conscious commitment to the belief that more is unknown than known and holding ones ignorance in mind as an intentional object, is the cognitive relation at the center of an epistemology of inquiry.

One last word of introduction. The inquisitive stance can be thought of as an explanation for the recursive interexchange between knowledge and ignorance mentioned in chapter 2 section 1. This basic relationship of exchange needs a mechanism of action. Ignorance does not transition into knowledge and knowledge back into ignorance by magic. There needs to be an underlying explanation. The key to this mystery is inquiry and its tools which are questions, that are like little hooks that latch onto the unknown and reel it in.

The process of inquiry does not begin with the question. The question is more a product of underlying mechanism. Indeed, the question is more toward the middle rather than the beginning of the process. The process begins with a metacognitive awareness of ignorance. To bother asking a question you must first realize there is something you don't know, and this something is definite enough that you can fit it into words. The awareness of this ignorance is what causes the question. The cause of inquiry, both its instigating factor and its final end goal, is

thus a peculiar type of intentional object or "mental picture." The object of inquiry has no positive signification or presence initially, the kind of mental content that would come to mind if one were answering a question rather than asking it. Instead the intentional object which grounds an inquiry is a privative, an absence; a discontinuity in knowledge that is noticeable precisely for the fact that it is nothing, a gap that is demarcated by the edges of limited knowledge not lining up. The inquirer must fix this intentional object of ignorance in mind. This is what grounds the process of inquiry, initiating it, and providing the reference point from which it evolves. From there on forward, the inquirer monitors how much this awareness of an absence diminishes as information is sought out that satisfies its incompleteness.

The saying "the absence of evidence is not evidence of absence" has a particular application here. Clearly some sort of negative evidence is required to come to the conclusion that one doesn't know something, but it is less clear what this evidence of absence consists of. I know with certainty that I cannot speak Chinese, but what exactly did I do to find out? It was an immediate and obvious conclusion, no kind of intensive review or thorough validation of my knowledge base was required to draw it. Presumably my evidence for knowing that I can't speak Chinese is found in the desire to speak it and then confirming that I am unable. Even then, my awareness of this ignorance is easier to grasp than that, as I make no attempt to speak it and do not need to try and fail to produce evidence that I can speak it to know I can't. Its absence from my knowledge is immediately apparent to me. Surely there are more ambiguous situations than confirming to myself that I don't know Chinese, lacunae within my knowledge that I am less conscious of and cannot even "demonstrate by failing to demonstrate" a lack of knowledge. Absence of knowledge is evidence of ignorance. The problem then becomes: how does one "gain purchase" into this ignorance, establish an adequate concept of it enough to ask a question that is a decisive step to overcome it?

Ignorance that can be questioned is of a more superficial kind than deep ignorance that is afflicted by *hypocognition*, a term from cognitive linguistics to describe the phenomenon of lacking words to describe a concept. Ignorance for

which I can pose questions must be peripheral and mutable. It can be put in perspective, I can label its elements. Furthermore knowing that we don't know something by itself does not provide grounds for knowing what questions to ask about it. After all, we have nothing to work with. It is a raw, uninterpretable and absolutely simple datum: I don't know. Little can be done with that on its own. When we are truly stranded and stumped on an issue, we are stuck not only on the problem of our ignorance about the target domain, but also the secondary ignorance of not knowing which questions to ask that will pave a road out of the quagmire. If we are truly ignorant then hypocognition should come into play and prevent any attempt to parse it into a directed question which points the way out. This poses a bit of a paradox. How can I know enough about what I don't know to ask the question that will lead to a correct answer and the resolution of my ignorance? This would seem to mean that I must at least be partially knowledgeable about the subject of questioning, otherwise I could not fit into words a good question to elicit the required information.

To be able to form a question suggests that the ignorance that stimulated it is somehow fragile or unstable, exposed to the possibility of knowledge, encompassed. Deep ignorance, ignorance of ignorance, the unknown unknown, is unquestionable. Not only is the content knowledge missing, but so is the judgement necessary to evaluate the meaning of this absence.

And it is precisely *judgement* which provides a way out of paradox. The question comes not from the ignorance of the content itself which is the first cause of it, and concerning which the question seeks an answer, but from the faculty of judgement that observes that ignorance. After all, nothing can come of nothing, a quantum of ignorance is dead weight and nothing can be done with it. Judgement, however, is a higher cognitive faculty which presides over and above the ignorance it evaluates. More must be said about how judgement relates to the intentionality of ignorance.

Ever since Brentano revived the concept from its medieval scholastic origins, a common assumption about intentionality is that it is "the hallmark of the mental" and that *all* mental states have this property of "aboutness." Furthermore it is assumed that this mental state refers to a content, which is to

say, the object of invention cannot be null, it must have some "suchness" about it. "What is characteristic of every mental activity is, as I believe I have shown, the reference to something as an object." (Brentano, F., 1874 [1911, 1973], Psychology from an Empirical Standpoint, London: Routledge and Kegan Paul.) However, while maintaining that a mental state must have an object, Brentano also insists, and seemingly risks a contradiction by doing so, that this object need not necessarily exist outside the mind. "The terminus of the so-called \[intentional] relation does not need to exist in reality at all." This so called intentional inexistence was meant to highlight the argument that the mental object held its own independent reality as an ontological fact. Whether or not it had a corresponding referent in the external world was optional. The idea of Bigfoot is its own fact, regardless of whether there is an external referent to which it attaches. Myths, dreams, fantasies, have some substantial reality even though they do not depend on mind-independent facts. Hence abstracts such irrational numbers can still be entertained as factual entities without being traceable to an external referent that anchors its meaning.

All things said, ignorance exists. Despite it being a privative, a negative concept, this does not jeopardize its ontological status. Its existence is purely relational: it exists within the "interstices", the cracks between knowledge, in the same way the space between two buildings exists by virtue of the distance between them, but would not be said to exist without them. Awareness of ignorance is thus the awareness of the "presence of absence." It is the discovery that something is missing.

For judgement to take effect Brentano also held, it requires a presentation on which to operate: there must be a representation serving as a logical content to which judgement applies. Taken together these claims present an interesting and unusual case when interpreting ignorance as an intentional object.

Whether or not the intentional object has an external referent, it ought to be effable. If one knows one doesn't know something, there should be a thing to which it corresponds. Thus the awareness of an object of ignorance ties the mind to it, establishes a relation or tether. Anything that is consciously apprehended should have enumerable properties that with sufficient effort might be listed.

Even the paradigm of a privative concept, mathematical zero, has attributes. Zero isn't a pure nothing about which no more can be said. It is even, neither positive nor negative, an integer, a multiple of all numbers, and when added or subtracted to a number it always yields the same number, etc. The same cannot be said for negative ignorance, except with a catch. The *possibility of discovery* of such properties is implicit in the awareness of its object.

Negative ignorance as an intentional object has no qualities. It has the *possibility* of such properties, and it is the pursuit of their discovery which motivates investigation through the search for an information source that satisfies it. Surely this absence of properties does not show when one is positively ignorant, delusional or mistaken, as a delusional or mistaken belief is no different from a correct belief in terms of the availability of conscious access. A false belief is about something just as much as a true belief. The problem, precisely defined, is that when we have the belief that we are ignorant, there is nothing to which the object of this belief can be related in the mind, since we cannot represent to ourselves what we do not know. Therefore, investigation can only proceed from negative ignorance, as positive ignorance starts from falsified content that misdirects and leads to *mis*information sources.

The complication regarding intentional ignorance is directly related to Russell's work on establishing grounds for the justification of true negative existential beliefs. How can one assert a negative truth claim (e.g. Santa Claus does not exist), if there is no referent to represent what the proposition is about? If the referent designated by the proper name does not exist, then the name "Santa Claus" has no object, thus the proposition is not about anything at all. If it is about something, that object must exist. Russell held that one must have acquaintance with- that is, sensory experience of- a particular individual object to assign a logical proper name to it and therefore to meaningfully denote it (Russell, 1905). To denote the meaning of the name "Santa Claus", it must be recast as a description of a concrete individual. For example, "a jolly, bearded, red suit wearing man who lives in the North Pole and flies around in a magical sleigh to deliver presents to well-behaved children on Christmas." Thusly imbued with a definite description, the world of empirical objects can then be

scanned to see if such an individual fitting that description exists or not. As such a definite description is then open to acquaintance while the proper name of a nonexistent object is not. It is impossible, so it is argued, to empirically establish the existence or nonexistence of what a proper name designates without a definite description of that object. Of course, the image and idea of Santa does exist, so the definite description of Santa as the "imaginary folkloric character who is jolly, bearded, wears a red suit…" does pick out an existent.

To tie this back to the problem of intentional ignorance, no such definite description is possible for a single itemized unit or "quantum" of ignorance such as "I don't know Chinese" in the sense that to describe what one doesn't know definitely is to then know it. The more you describe what you don't know the more you know about it, the more the object of ignorance to which your thought is fixed diminishes in proportion to which the object of knowledge grows. One cannot describe what you don't know in any detail without coming to know it.

The confusion about how one can be aware of what one doesn't know subsides once such self-recognized ignorance is understood to be *metonymic*. It is a kind of placeholder, and a placeholder is designated with the expectation that there will be something at some point to fill it. To recognize one's ignorance is to substitute the empty privative (the absent which is present) for a kind of algebraic variable that stands for an as of yet unspecified definite description of an object of knowledge. It is to become aware of one's potential for knowledge. It's like saying "**something* is going on outside", "*someone* will stop by to deliver the mail", or "do *whatever* it takes." These statements substitute an unidentified particular object with a generic term (something, someone, whatever) that functions as a kind of logical filler, expressing that there is in fact a referent that is predicated by the rest of the sentence even though it cannot be specified a proper name. These incomplete descriptions, if substituted for a definite description ("a parade is going on outside") would amount to complete knowledge, a known known, rather than a merely known unknown.

We do not have a clear idea of the intentional object of ignorance, because by definition we not know anything about the referent to which it is connected other than that it is unrealized potential knowledge of some subject. Otherwise, I

would have to be acquainted with Chinese to know I don't know it, which is a contradiction. But self-recognized ignorance is always about something, there is some relation between the mental state that recognizes it and a corresponding unrealized potential. (Or at least, this connection between awareness and the potential for knowledge is the precondition for investigation; as one may have an awareness of unobtainable knowledge, for which any investigation in pursuit of would be futile.) If one hears a loud bang outside one does not need to know what caused it for the hidden cause to be an intentional object that the awareness is about. Clearly in this example such a hidden cause does not need to be effable at the moment of recognition for it to be the object of intention. By this metonymic means, judgement can "gain purchase" on a discrete "portion" of ignorance and hold it in mind as a distinct object. To raise ignorance into conscious awareness is to assign it a metonymic label which can then be fitted into propositions alongside positive knowledge.

All of this is admittedly a bit complicated, but it is necessary dirty work that must be sorted out in order to progress to a more well-founded explanation of the inquisitive process. All in all once ignorance has become the object of intention, deliberate, directed learning becomes a live option. Without such intentionality, learning can only be incidental and passive, one cannot say to oneself, "I know what I must learn." The argument is that such goal driven learning cannot be engaged without first intentionally objectifying the discrete ignorance connected to it. The hole must be found before it can be filled. Therefore to fix one's ignorance in mind is the first necessary step of proactive investigation.

On one reading, deliberate learning consists of two components that can be denoted by their respective inward and outward dispositions. The intrinsic component is this intentional objectification of ignorance described above. This discrete recognition of what one doesn't know results in indexing it as a bounded finitude. It is to break off, in other words, a piece of quantified ignorance, carved at the joints, and detached from a vague and undifferentiated mass of deep ignorance about which there can be no definite awareness. Now, to mobilize learning, what can help it along is what has been called epistemic emotions such

as curiosity or wonder which serve as motivators to compel the search for knowledge.

This brings us to the extrinsic component of deliberate learning. It is the process of active sensory exploration and searching for information sources within empirical space. In order to trigger the subjective conditions by which one judges one's ignorance to have been eliminated, some criteria of adequacy for external information uptake must be met. What are these criteria of satisfaction? The answer to the question must come from outside, and one must be reasonably confident that the answer one receives is sufficient. This raises the question, how does one know one has learned enough? If, as we have assumed, the awareness of one's ignorance is what compels directed inquiry, and if we assume that inquiry ceases once this awareness has dissipated because there is no more ignorance of of which to be conscious, how can one know one has learned all there is to know without presuming that knowledge (Meno's paradox)? This sort of subjective halting problem has come up elsewhere, but Meno's paradox is perhaps the most concrete example of it. There is no way to know before hand that the external information source is adequate, as this would presume knowledge it would be impossible to have. And there is no way to know has learned all that is necessary to complete a task, at least not until one attempts to accomplish it on the basis of received information and succeeds.

Even though one starts not knowing how to align the object of ignorance with the appropriate information source, the line must be drawn somewhere. You must start without necessarily knowing where to start, selecting instead a vague general area guided by unarticulated hints in a particular direction. Just as how "left or right" divides the world and two, a series of successive binary approximations will eventually narrow the whole universe down to a pinpoint. There are a few factors at play in determining what must be investigated and how much: the depth and complexity of the sought for domain knowledge, the appropriateness of available information sources, and the precision with which the leading questions are defined. Now to address these in turn.

A complex subject obviously requires more cognitive effort to unpack. The extent of its complexity may preclude asking the right questions at least until one

has learned more about it through a kind of undirected and loose preliminary study of it. Its complexity may also increase the uncertainty that there is still more to find out, causing the opposite problem of when to end the investigation. A complex question is one that does not have a yes or no answer. There is in fact, uncertainty at both sides, at the start and at the end of the investigation and the appear to be related by the same problem. At the start there is uncertainty about whether one is asking the right questions or barking up the wrong tree. At the end, there is uncertainty as to whether one has exhausted the information source. The only certainty one can have is that one's inner sense of ignorance, quantified and identified by one's questions, has run out because a subjective sense of adequation has been received, analogous to a hunger that has been satiated. Such a feeling of informational satisfaction is of course not an objective indication that everything there is to learn has been learned, only that the agent's own curiosity has reached its end.

One solution to problem of determining the adequacy of information sources is that we rely on education and the authority of more informed people to guide learning. Someone else draws the line for us. Using a curriculum, an arranged and prepared collection of information sources, cuts out the guesswork. Such shortcuts can evade the often aimless "random walk" behavior of an unsystematized and unorganized study, a study that is lacking the right questions and therefore struggles to target the most useful unknowns, the discovery of which would simplify and facilitate learning the rest. A simple example of this should suffice to demonstrate the principle. A child who wants to learn multiplication without first knowing they must learn addition will have a harder time, as multiplication is just the expedited form of addition of a number to itself as many times as the multiplier. For this reason addition is always taught before multiplication. In many domains of knowledge there are such "fundamentals" the learning of which accelerates learning the rest. Someone who is unable to mount a successful investigation is someone who is unable to identify which key unknowns they must direct their attention toward that unlock the potential to learn the rest.

Finally, a well-defined question clearly isolates the object of ignorance. Such

that by comparing each term of the question and the answer proposition, it should be possible to tell by the form of the answer alone whether it satisfies the question. To use a trivial example one can intuitively know the question, "where is Mary and what is she doing?" is not satisfied by the answer "she is in the office." The response in this case obviously leaves the condition "and what was she doing?" dangling. After all the question is really a conjunction of two questions, "where was Mary" and "what is she doing" each of which must have an answer. Similarly, the response "anywhere but here," is inadequate. It fails to be disambiguate. Even "maybe she is in the office," fails to completely satisfy the requirements, because it expresses doubt on the part of the speaker. If an answer is to be accepted unconditionally, it must not add uncertainty to the questioner's uncertainty.

This last note about the formal symmetry between the terms of the question and the answer does not tell us about the accuracy of the semantic contents of the answer, but it does provide a basic analytic a priori check on the adequacy of the answer based on how their terms match up. Even "she is at the office" is a synthetic proposition that cannot be derived from the sentence alone and is only confirmable by experience. But the response that she is doing "anything" is a non-answer a priori, she is doing whatever one thing she is doing. A similar flaw would appear if the response was " she is not at home and not in the office, but perhaps she is at the other places she usually goes." If the answer says too little or too much it's inadequate. The answer must fit the conditions of the question. If the answer does not reduce ambiguity but merely reproduces it, it is inadequate. Vagueness is unacceptable: *the answer must provide actionable information*. Like the notorious Glomar response made famous by the CIA, it "neither confirms nor denies", failing to give actionable information by which the questioner can further narrow and target their information search. It's therefore self-evident that an undersupplied or oversupplied answer that does not align with the logical form of the question must be insufficient. On the contrary a modal question, for example Q: "what might Mary be doing" does afford a modal disjunctive response, A: "She might be doing x or y or z." Again, the point still holds that the question constrains the answer.

Similarly the complexity of the answer ought to be proportional to some degree to the question. This is somewhat analogous to Carl Sagan's famous quote, "extraordinary claims require extraordinary evidence." While it is certainly possible to resolve a detailed question with a simple yes or no, the probability of being able to do so diminishes as the complexity of the question and its answer criteria grows. This does not mean that a complex question is so because it is verbose, even a simple adverb can imply complexity: "is it raining out my window?" affords a simple yes or no answer, while "why is it raining…" demands a lesson in meteorology.

Yes or no questions are an important class of questions. They are recognized as a distinct type for good reason. What makes them special? It is that they are *purely logical* in that they exclusively accept the binary truth values of classical logic, true or false. They are the epitome of precisely defined. Such questions are strictly logically constrained to the fullest, forcing the answer to confine itself to an absolute commitment to one of two exclusive truth claims. Must every question reduce to simpler questions that accept a yes or no answer, as an information theorist might hope? A bit can be defined as a yes or no (digital) answer to an absolutely simple question, such as "is the lightbulb on". The game of 21 questions can be used to narrow down and identify an object from the universe of all possible objects. This in itself is a paradigm example of a structured inquiry interrogating an information source and narrowing down the possibilities. I am unsure if in principle all answers can be reduced to a combination of yes or no atomic propositions, although this is certainly food for thought. As it turns out Hintikka (2007) suggests that perhaps every complex question can in fact be re-described in terms of yes or no questions, however this would entail an inefficient increase in verbiage as a general question is decomposed into a series each of which must be asked separately. In a class of 21 students, one day an apple on the teacher's table is stolen. The teacher can ask the whole class, "who stole the apple?" They could also ask each student individually "did you steal the apple (yes or no)?" While the precise binary nature of yes or no questions makes them appealing units for an imagined rigorous formal calculus of inquiry, it smacks to me more like wishful thinking

that language would submit to such simplification without becoming distorted.

Moving on, to return to the problem of adequacy, it is clarified further if we take a more humanized view of the investigation process and see how emotions play into it. The process is not a strictly "cold" logical and informational cognitive operation. As Helen de Cruz states, feelings such as awe and curiosity are epistemically productive. We do not require a formal definition of adequacy if we persist in being curious about a topic, in the words, if the investigation is no longer guided by definite questions or hypotheses but rather by a limitless desire to find out everything about a subject. Curiosity is learning for learning's sake. Wonder orients the agent toward open-ended discovery and gives a supra-rational reason to keep pushing for knowledge until the exhaustion of the emotional impulse. Emotions are not merely passive states that come over us, they are calls to action \([https://aeon.co/essays/how-awe-drives-scientists-to-make-a-leap-into-the-unknown]). Awe and curiosity permit the mind to open to new possibilities. Cruz calls this opening up of the mind cognitive accommodation. "The need for cognitive accommodation makes you aware that there is a lot you don't know…it is also an epistemic emotion, because it makes us aware of gaps in our knowledge." To consciously orient oneself to the unknown and to suppress the urge to fall back on preconceptions is to relax and soften up hardened convictions and abandon illusions of knowledge completeness.

Psychologist Daniel Berlyne is noted for associating curiosity with a mood of deprivation and original work on the psychology of questioning. Uncertainty is often-but not always-unpleasant, and the recognition of it spurs a drive to eliminate it, which involves curiosity. Indeed several of his ideas anticipate facets of the intentional stance. He harkens back to the theories of Gestalt psychology and its "principle of closure" in which the mind "fills in the blanks" of an incomplete representation. Curiosity is uncontroversially what drives the impulse to close that gap. The idea in either case is that there is a gap, that it must be closed, and the question is how to go about it. In thee 1954 paper *A Theory of Human Curiosity* Berlyne reviews the science of his day on the topic and settles on a conflict-driven theory of curiosity. Centering his theory on behaviorist

analyses, questions are viewed as self-stimulating verbal behavior and the goal of curiosity is to remove an unpleasant stimulus. The centerpiece of his theory is that puzzling or unusual stimuli act as irritants the source of which is then actively explored.

(Footnote: Tangentially there are a few gripes I have with this work. Although the intentional object of ignorance is implicit in this "gap theory" the assumption is not that there is a gap within the representation, but nothing to start with. It is no filling in the blanks. The whole page is blank. The riddle is how novelty is overcome. Additionally, behaviorist conceptions of questioning, called thematic probes in the jargon, are considered to be the initiators of the exploratory behavior, which is self evidently absurd.)

([https://static1.squarespace.com/static/53a79084e4b01786c921de45/t/53a86486e4b009ec07711b59/1403544710847/A+Theory+of+Human+Curiosity+(Berlyne,+1954).pdf]). Berlyne also distinguished between epistemic curiosity—the urge to know ideas or abstractions— and perceptual curiosity, the desire to discern what is going on in the sensory environment. As it happens, more recent research suggests that epistemic curiosity its more associated with a positive sense of wonder while perceptual curiosity negative affect. ([https://www.psychologytoday.com/us/blog/why/201706/why-do-we-ask-why]) On an animal level this makes some sense. The realm of abstract concepts may be confusing at times, but it can't hurt us, and it's hard for anybody not to find something that isn't immediately within range of the senses that worth's wondering about. Unknowns stirring in the physical environment however could be something dangerous or at least annoying.

Also intriguing is brain imaging studies which have associated curiosity with the nucleus accumbens, a structure strongly linked in reward and motivation. The nucleus accumbens is known to greatly involve dopamine in its mechanisms. It receives projections from dopamine-producing neurons, and the excitatory activations this induces is thought to play a significant role in experiences of reward, pleasure and motivation. ([https://pubmed.ncbi.nlm.nih.gov/25284006/])

The reason for mentioning this brain research is how the humble molecule

dopamine and its role in the nucleus accumbent can explain so much. Dopamine is commonly referred to as the "pleasure molecule", implying that its expression is the biochemical source of pleasure, but this is misleading. Dopamine neurons not only respond when the animal receives a reward but also when a stimulus… predicts a reward. Such reward-predicting stimuli are conditioned rewards and have similar effects on learning and approach behavior to real rewards (Schultz, 2016). There is a sense then that dopamine is just as equally the curiosity molecule, or in the most general terms, the probability molecule. It has an essential role in orienting the brain toward stimuli which may be promising. In doing so it inherently encodes a probabilistic guess which carries an expectation of a possible reward. Precisely such an element would be required to orient the brain toward the pursuit of the unknown, especially if there is a payoff.

While the scientific research on curiosity and questioning is deep and interesting, my personal preference is that this is a subject better handled by philosophy. Specifically when it comes to questioning. Much is to be learned from the psychology of questions, but understanding them on the most fundamental level benefits from strictly viewing them as logical constructs.

These converging tendencies of self-recognized ignorance, the intentional objectification of ignorance and epistemic emotions, terminates in the decisive act of posing a question. Philosopher Lani Watson has made a study out of questions. Watson is concerned with questioning questions. Such meta-questioning unveils a ripe space for thinking about thinking that has been surprisingly overlooked in the history of philosophy. If as is often said history repeats, or as the original Santayana quote reads, "those who cannot remember the past are doomed to repeat it," it can only mean that one has failed to step outside the first-order stream of events and neglects to adopt a second-order, distanced and removed perspective on it. Socrates, Watson points out, embodied the questioning spirit of philosophy through the Socratic method but declined to question this method of questioning itself. Fruitful recent work on the logic of questions, such as Hintikka (2007), tends to restrict itself to a logico-propositional analysis of questions as logical modality rather than as a more comprehensive cognitive psychologial act. Watson goes beyond the propositional form of a

question and reveals layers of subtlety in the act of questioning as a form of engagement with the world.

There are several findings to this end. A question, doesn't necessarily need a one-to-one answer. It may prompt many possibilities which, even in spite of their variety and in spite of no method for selecting the true answer from these options, still yields further information for further questioning. (Arguably this is not inconsistent with the previous discussion about the adequacy of answers, because the fact that an answer opens up many more possibilities for questioning means that it was inadequate.) What's more, Watson points out that a question does not require explicit linguistic expression and can be any information seeking intentional state and need not be verbalized. A google search query is an interesting case. One must simply have an unresolved inner state of uncertainty and seek to tap into an external information source that would reduce it. Of course verbal queries are more readily analyzable because of their formal, symbolic properties, but any information seeking behavior qualifies as a question broadly realized. By getting out of bed and opening up the window blinds to see if it is raining, one's behavior implicitly embodies an interrogative attitude without speech. As Watson asserts, "a question is an act." (Watson, Loni, What is a Question? The Philosophers Magazine, 82 (2018)). By extension, not every answer has to be in a verbal form, it can be any type of confirmatory sense data. You do not need to rely on someone else to tell you what the weather is like, you can look outside.

In summary, all inquiry begins with the intentional recognition of one's ignorance as a metonymic psuedo-object. A question or system of questions existentially quantifies or "draws a circle" so to speak around a portion of ignorance ("there is some x I do not know"), specifying it so that an actionable search for information sources can ensue. Epistemic emotions help spur this process on when the means for accurately evaluating the true state of one's ignorance—information incompleteness-relative to the complete objective information is lacking. These emotions are the experience of fascination: you keep learning without searching for answers. Investigation continues until the terms of the answer are judged adequate, and a subjective sense of parity or

equilibrium of mutual information is achieved between the mind and the information source, that is, the information source is exhausted and the ignorance satisfied, like a cup of water being poured into another. Inquiry proceeds to the extent that the decrease of ignorance is inversely proportional to the increase of questions answered until either one's awareness of one's ignorance or the information source dries up and no more can be found. Once one is no longer cognizant of one's ignorance as an intentional object, a feeling of satisfaction is induced and the inquiry ceases. This feeling of satisfaction is not necessarily accurate, however, because one may not be aware of the true extent of one's ignorance relative to what remains to learn. Deep ignorance-unknown unknowns—may still remain with respect to the topic after one has eliminated the intentional apprehension of surface ignorance-known unknowns.

This entire process can be illustrated by diagram:

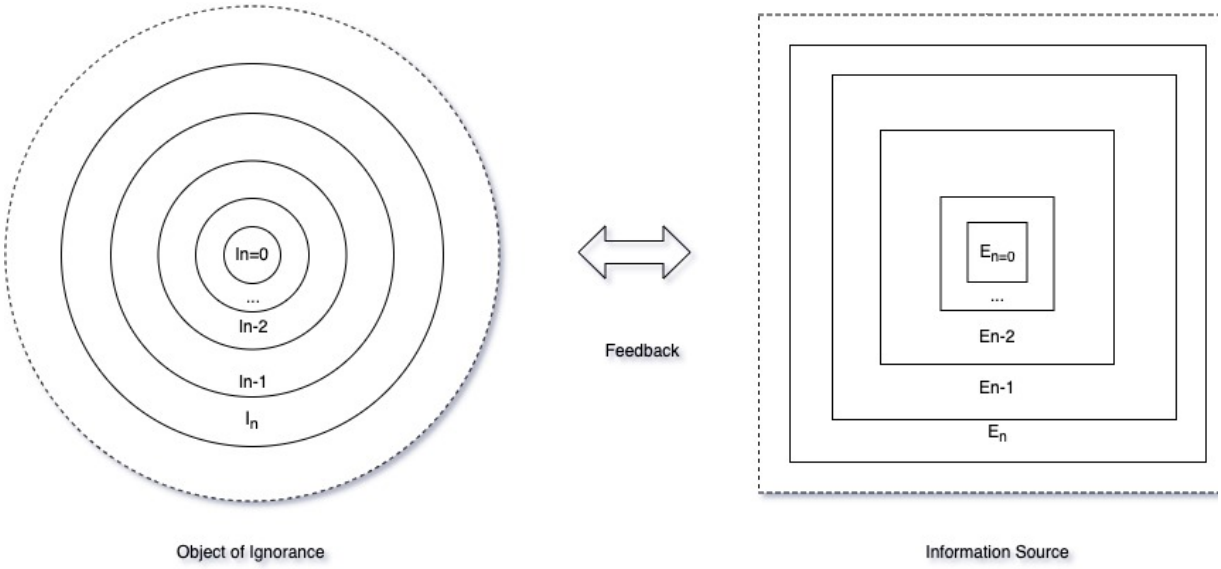

Here is effectively the mapping of the codomain (image) of the function f: I → E. Where each element of I is a state of ignorance that corresponds to an element of E which stands for an information state of an external information source. The dotted line surrounding both sets I and E represents the "horizon of uncertainty" which envelops both domains. This is meant to represent that the cardinality of either set is undefined, and that true context of I or E may be greater than n. In other words there may be more external information En+y than what In selects as the initial baseline cutoff point. One's awareness of one's ignorance, may be

less than the full extent. There may be more ignorance In+y than is encircled by the investigation. Thus, the sets are bordered by a fuzzy boundary. One cannot know the extent of one's ignorance and also therefore the extent of external information that is fit to answer to it. There must however be an initial state f(n) which initializes the function, a leading question as it were. The investigation begins within the horizon of uncertainty wherever a definite object of ignorance In can be identified. The double arrow stands for the bijective relation and feedback that occurs between both domains as n approaches 0. Like layers of an onion, the object of ignorance and the information source are peeled back until there is no more ignorance to be subdivided into more granular investigation states and no more external information to be unpacked. Each layer can be thought of a round of questioning and answering that successively iterates to the next step.

An answer to the question is simple if the question is simple. The question of whether it is raining outside can be confirmed by a simple yes or no answer by looking out the window. In this case, the depth of the inquiry function is f(0). Much more complex questions such as "what were the causes of World War One?" probably have a much greater depth.

This model of the relationship between the object of inquiry and an information source is of course idealized and simplified. However even where it fails to capture the true complexity of inquiry it can still be instructive. The spatial metaphor which views uncertainty as "atmospheric" and enveloping the domains is imperfect. Uncertainty does not simply exist on the outside, but also penetrates the process of learning in the course of its unfolding. One can imagine "ghost layers", in the information source, missing pages as it were, that halt or sidetrack the process. Imagine for instance a layer for Iy which has no corresponding layer Ey, and is instead replaced by another dotted line rectangle denoting a null layer of uncertainty. One may reach a given state in the function, where the state of I has no corresponding state of E. There is no reason to assume that the process must be linear, no reason to assume every question must have an answer and one and only one answer, or that an information source is perfectly adequate to accomodate the inquiry. An investigation doesn't necessarily stop if

in the course of the iterations of this cycle it encounters missing external information. Perhaps the solution is a recursive subroutine ("Why is that information missing?") which occurs in such a case. The missing information spurs a sub-investigation and creates a new object of inquiry, and the process is executed there until the missing information is returned before continuing to the higher level inquiry. In all the main idea is that there are coordinated relations between both the questions asked and the answers received that progress in step until the object of ignorance is completely eliminated while in proportion the information source is exhausted.

Inquiry is also readily diagrammable as a more familiar flowchart (expert system). A flowchart better illustrates the the conditional (if/then) logic of question and answering rather than a proportional extractive relationship between an information source and an inquisitive agent under ignorance. Say for example someone wants to know whether it is raining or snowing outside. Schematically they would proceed in this way:

Is it cloudy outside? (Y/N):

If N then : end query.

If Y then: is it warmer than 32 degrees?

———- If Y then: is it raining?

—- If Y then: end query

—If N then : is it snowing?

and so on…

Obviously, this toy example is not representative of how a person would actually approach checking the weather. They would skip to the chase and ask if it is raining or snowing, bypassing the assumed prerequisite of "is it cloudy." Or they would abstract the question of whether it was raining or snowing by asking what the weather is in general, which would give them the maximum of information with the most efficiency because they could also learn as a free bonus whether it was sunny or just cloudy without precipitation all in one question. Or they'll just look outside and see for themselves. In reality not all questions are true or false, which further complicates matters. The fact that answers can have deeper semantics than yes or no vastly expands the amount of

information that can be both expressed by and elicited by a given question. Even so, the logic of a flowchart captures the relational, treelike conditional reasoning of questioning.

For more complex investigations this formal structure of conditional dependencies is more relevant. A police detective engaged in interviewing a witness has no such heuristics. They cannot assume anything. They must start from the most general of questions and continue to specifics, trimming off branches of conditional possibility until they parse the one truth.

Here it gets more complicated. The conditional logic of questioning is not such a simple, stochastic process by which the next question to ask is dependent only on the current state and the previous answer. What must also be taken into consideration is the summation of prior answers everywhere in the tree. The simplified model assumes a perfect information source, which always gives the right answer at every state. If the information source has a probability of being wrong, however, the asker must factor in the posterior probabilities of their belief based on how well the answers to other questions are likely to infer it. (And of course it is a bit of a fantasy to suppose there are any information sources without a probability of being wrong.) Otherwise, they may end up asking the wrong questions based on wrong answers, which is a sure sign that an investigation has been derailed and is headed toward failure and the murky depths of positive ignorance. The goal of an investigation is truth, but the result may just as well end in falsehood. The tendency to view these phenomena in simplified, logically abstracted, skeletonized terms leads to idealizations that tend to ignore gritty practical realities. The possibility of being wrong or misled during investigation must be appreciated.

To validate any given answer for consistency, the truth values of the whole chain of questions everywhere in the tree must be summed together in a fashion reminiscent of Bayesian belief models in which the weights of various belief states are readjusted depending on how much the actual input matches the expected value. In other words, one answer may fit the local requirements of its particular pathway down the tree, but it may conflict with an answer on an entirely different branch, so all data must be correlated. Granted, there are critics

of this Bayesian modeling of question-asking, notably Popper and Hintikka, but as we will see Hintikka's critiques in specific are not strong enough to warrant discarding Bayesian inference as a useful model of probability expectations that is generalizable to inquisitive behavior.

To illustrate the point, consider someone who is caught in a lie. A liar must engage in a bit of double entry bookkeeping, monitoring both the reality and the fiction. To fish out the lie, an interrogator can ask questions that trap the liar into a logical corner of self-contradiction. A liar can keep up the facade so long as nothing else that they talk about, even seemingly tangential details, conflicts with the version of events to which they attest (or the all important evidence, which we will factor in soon). A crime suspect trying to save their own skin might lie in response to a given series of questions that has a specific assumption in mind. They may succeed to evade detection down this particular question path. Yet another chain of questions, indirectly related to the previous, can expose an inconsistency on the overall picture. The most explicit example of this dynamic is to separately interrogate both a suspect and a witness and to see how their narratives align. The witness, who is presumably neutral, can serve as a reference while the suspect's words are matched against theirs. Indeed such techniques for cornering a deceptive suspect into a contradiction are standard practice for detectives:

> Some experienced criminals or persons who have committed well-planned crimes believe that they can offer an alternate explanation for their involvement in the criminal event that will exonerate them as a suspect. An investigator may draw answers from this type of suspect by offering the same proposition that is offered for exoneration. This is the opportunity for a suspect to offer an alibi or a denial of the crime and an alternate explanation or exonerating evidence. It can be very difficult for a suspect to properly explain away all the evidence. Looking at the progression of the event, an interrogator can sometimes ask for additional details that the suspect cannot explain. The truth is easier to tell because it happened, and the facts will line up. In contrast, a lie frequently requires additional lies to support the untrue statement. Examining a statement that is believed to be untrue, an interrogator can sometimes ask questions that expose

the lies behind the original lie.

(Gehl, Rod & Plecas, Darryl. (2016). Introduction to Criminal Investigation: Processes, Practices and Thinking, chapter 8. New Westminster, BC: Justice Institute of British Columbia.)

The discussion up until now assumes that the interrogation process is entirely self-contained within the inner logic of the linguistic conversation. In practice, evidence is the decisive factor. Evidence can be thought of as what defines the grounds for belief of the investigator's question system that assign it satisfaction criteria. A shrewd investigator will base their questions on the evidence and they will judge the answers on the basis of the evidence, in conjunction with the answer system's logical self-consistency. These interactions between evidence and the verbal exchange together constitute an expanded framework of inquiry and are by no means independent. All told, the story the suspect tells must at first not contradict itself before it contradicts the evidence. Just like an alibi, evidence can be incomplete, fabricated, or misleading, and it is possible for a good liar to work around it. However, if they hoist themselves by their own petard it undermines the integrity of their rendition of events just as impactfully. And if they can't keep their story straight then how what they say compares to the evidence isn't even required to cast doubt.

The flowchart illustration sketched above is essentially the same as the concept of forward chaining, or in its more modern iteration, depth first search, used in some artificial intelligence models (Russell & Norvig 2009, p. 337). Originally, forward chaining was used in early AIs called expert systems. These programs encoded if-then representations between entities and were then fed input, which the system checked against its knowledge base. A medical expert system for example could be fed a list of symptoms and look up a diagnosis match in its knowledge base. Researchers found that these implementations required excessive amounts of hardcoded, preinstalled data about relational entities, which effectively meant that their models were static, unscalable, and had no mechanism for learning. The world is constantly changing and new facts

are discovered and old ones reevaluated. It would be impractical to continually manually update the expert system so that its knowledge base grew in step.

This fault of static representations in expert systems is related to the problem of meta-ignorance. From within a system of knowledge, how is it possible to judge that it is comprehensive? Such knowledge will only permit you to understand what it allows you to understand, enforcing a blindness on all else. How is it possible to know that a system of questions is sufficient or has been answered completely? Or alternatively, that a knowledge base is adequate to solve a novel problem? It is the riddle of epistemic incompleteness. As a Japanese proverb goes: a frog in the well knows not the ocean. How to know there isn't ever one more question to be asked, that no stones have been left unturned? Certainly one's background knowledge place limits on what questions can be broached. At a certain point a wall might be hit, at which time it may be not that everything all facts have been exhausted, but that one is mistaking the wall for the edge of the world. And when we are in this bind, it's hard to know which of these is the case because the information available, or unavailable as the case may be, applies equally to both. As Schopenhauer said, "every man takes the limits of his own field of vision for the limits of the world. This is an error of the intellect as inevitable as that error of the eye which lets us fancy that on the horizon heaven and earth meet." Such self-limitation means that while an inquiry will draw a line and stop, where that line stops might not coincide with the limit of what's knowable. Even the full mobilization of the powers of inquiry to conquer ignorance is ultimately swamped by it. In the long run entropy always prevails. How to avoid being the frog in the well, oblivious to the ocean?

The missing piece here and the key to defend against mistaking the well for the ocean is to expect that there is always at least one more stone left unturned. If there is indeed one, you will be more prepared for it, and if there isn't, there is no cost to this falsehood because you already knew everything anyway. Leave every situation suspecting that you overlooked something. Of course one may be blind to one's own ignorance, and this pure ignorance, without any connection to knowledge to surface it from the abyss, may never come awareness. One may falsely believe an inquiry to be settled once curiosity has been satisfied or all

questions one can think of are exhausted. It's better to always bet that there is still probably something you don't know even if not for any other reason but for the sheer mathematical fact that more is unknown than known in general. This fact can be used as a statistical base rate to which one can always refer to as an invariable guiding principle. (Indeed as was demonstrated in the previous section on computational limits, it is provable that the quantity of ignorance always exceeds that of knowledge.) It is revealing that the way forward for AI beyond expert systems was through the development of machine *learning*, as it is abundantly clear that learning is the ultimate adaptation to ignorance. A system starting off with minimal knowledge but a robust capacity to learn will eventually exceed a system, such as an expert system, which is preinstalled with a solid chunk of knowledge but no way to add to or revise it. This tacit incorporation of ignorance management into AI models has helped make them drastically more powerful in recent decades. In line with everything argued herein, these ML models suggest it's not what you know. It's what you can know, the openness to the universe of the unknown, which makes all the difference.

The solution to meta-ignorance is meta-cognition: the system of questions in turn can be subjected to questioning. One can ask oneself, "am I on the right track? Am I missing something?" As Peirce noted, beliefs are fixed— and thinking stops—when a certain level of doubt has been sated. On the flip side, this can be reversed, belief can dissolve back into doubt. In *The Fixation of Belief* Peirce takes a more unidirectional view of belief formation, assuming that doubt, through thinking, typically passes into belief but belief not back into doubt. People usually have a strong dictate for doubt and crave self-certainty. Yet doubt never loses its utility in prompting beliefs to be reevaluated and thus to enable learning. Such *second guessing* is not a sign of intellectual insecurity but a strategy to foolproof assumptions and catch errors.

When all these threads converge into one it becomes the inquisitive stance. In summary, the inquisitive stance is the methodology by which a cognitive agent orients itself toward its own ignorance and renders it as a discrete object . Initiated by the intentional objectification of ignorance, driven by epistemic emotions and assisted by sensory exploration; directed by targeted questions and

confirmed by answers, the agent seeks information which counterbalances this initial deficit. A kind of work cycle is in play, in which information flows across a gradient, from the data rich outside world into the ignorant mind. It's a deceptively powerful statement to say, "I don't know *that*." The living possibility, or at least the enhanced likelihood, of obtaining that knowledge is a direct consequence of this realization. Without consciousness of one's ignorance as a distinct object- although an empty, null object, the presence of absence- the only other way to discover what one doesn't know is by accident or caprice. Unless a deliberate project to banish recognized ignorance is undertaken, the diminishment of it is contingent on passive experience and incidental encounters.

Often the pursuit of knowledge is held up as one of the most noble aspirations of which humanity is capable. Yet this is not the right way to look at it. The knowledge is not "out there" to be pursued. At best, only information can be said to exist externally, but information is not knowledge. knowledge exists only in the mind, and it does not exist until the mind searching for it has derived it. So what exactly is being pursued? You can't find what cannot be found. The same goes for the term "knowledge acquisition", although for the sake of convention I still use it. What is there to acquire? The conclusion comes before the premise. This is circular.

It is more meaningful to speak of fleeing from ignorance. The goal is not to pursue knowledge but to escape being ignorant, and these are different concepts. By attending to ignorance we necessarily hone our conception of it, highlight its outlines, identify the shape of the missing puzzle piece from the arrangement of all the surrounding pieces. We find out what we need to move away from. Information may be the prerequisite of knowledge, but all the information in the world does make a difference unless the keyhole in which it fits is at first located. It's more accurate and insightful to say we flee ignorance than pursue knowledge. It is by getting away from our ignorance, of which we have become so aggravatingly conscious, that we thrust ourselves toward information which is adequate to resolve it, and as this information is absorbed it is constructed into knowledge. As has hopefully been argued to satisfaction, this inversion is not a

mere tautology, a way of saying the same thing in double-negated terms. Ignorance is the starting point, the predominating epistemic condition, and knowledge is a mere reaction to it. knowledge is never as certain as ignorance. If nothing else we can be certain of some form of ignorance. knowledge in contrast can be fragile, reversible; it can expire, degrade, and be superseded. Contingent, conditional knowledge may be falsified as soon as the conditions on which it depends change. Not so ignorance conceived in its totality. Indeed it is much easier to definitely define ignorance than it is knowledge, which may explain why the latter has remained an ongoing challenge. Perhaps this suggests something about knowledge's peripheral nature. Is knowledge so hard to define because it is inherently mutable, always shadowed, bordered, infused and led on by ignorance? What passed for knowledge in the now outdated science of the 1800s is no longer considered knowledge, but who are we to conclude that the present state of knowledge is on a more solid footing?

Meanwhile ignorance remains. While some of its territory has been conquered, it seems to grow an acre for every acre claimed. If we allow that knowledge as I believe we should that does not have to be perfect and immutable, that it can be provisional, logically fuzzy, and pockmarked by ignorance, then what we are simply doing is crediting ignorance its proper place in epistemology by subordinating knowledge to it. As it happens, this is also the route to intellectual humility, the understanding that the world is more complex than we think it is. To the extent that knowledge can be said to be possessed, it is in that it reorders and clarifies the knower's relationship with their own ignorance, allows them to adapt. But just as size of any countable subsequence of an infinite uncountable series, no matter how large, is always infinitely less than the size of the uncountable series, so too is knowledge subsumed by ignorance.

Being asked questions is important to the theory of inquiry too. A question which cannot be answered is also a question which cannot be understood. An answerable question assumes a parity between the known and the unknown, it assumes knowability. An unanswerable question, even if it is well formed, still yields some information in the same way that null results do in science. It still points to something, and raises the further question about why exactly this

question can't be answered. It shows us dead ends and therefore helps reveal which paths remain viable. (Such dead ends can be "bracketed" as pending and uncertain and unbracketed later when more information comes to light, as they may only be unanswerable given the present state of inquiry.) By going unanswered a stumping question tells us what we can't know, it provides proof of inadequacy which again furthers investigation. It validates the integrity of the knowledge base. An unanswerable well-formed question raises the specter of structural ignorance, which is reflected not by simple lack of knowledge but by the limitations of the framework of existing knowledge. knowledge which prioritizes certain focuses and so neglects others and is inherently circumscribed. Newtonian physics is certainly knowledge, but it still falls short of the fuller picture, and incorporating relativistic principles into it —incompatible knowledge forms— unravels it.

For these reasons it is the questioners, who are by extension always learners rather than lazily self-satisfied knowers, who are the real epistemic powerhouses. They do not stop at an illusory boundary of falsely complete knowledge, the telltale sign of epistemic senescence, but continually push the envelope into the interminable unknown. They reside on the edge, continually revolving, turning over their minds, adapting their ideas to new discoveries. They anticipate complexity, and if they overshoot their estimation of it, they are still better off than if they underestimated it. They situate their thinking at the point of intersection between the edge of knowledge and the beginning of the unknown, at that one vital point where the two curves overlap. The best way to demystify these ancient philosophical enigmas about how we know is exactly this inquisitive stance in which the unknown, punctured by inquiry, continually flows into the known by having its information content siphoned off into it by the asking of penetrating questions.

Hintikka's Socratic Epistemology

Questions are doubtlessly an essential feature of the learning and thinking

process. Interrogative sentences were used as far back as the Epic of Gilgamesh in 21BCE, and while it can't be proven it's a safe bet that for as long as there has been language there has been interrogative speech acts. In all likelihood, curiosity as the primal emotive precursor of verbal questioning predates humans, and perhaps even primates. Chimpanzees exhibit behavior that can be unambiguously interpreted as expressions of curiosity. Chimpanzees who have been exposed to the sight of fire for the first time have been observed to display investigative behavior such as fixating their attention on it and using tools like sticks to poke and prod it (Byrne, 2013). Ants, in their restless drive to scan the area round the colony for food sources may not be exhibiting conscious curiosity, but they are behaving that way for the same reasons that curiosity exists in primates.

As fundamental epistemic constructs, questions and the process of asking them has understandably been the primary stimulus behind the history of ideas. After all you can have no development of ideas unless there is something to trigger change and the incorporation of new information. And rather than to wait to be proven wrong one can actively question and take control of one's ignorance. Such changes, if not always preceded by a question, in every case in which it is always stems from a deliberate engagement with the unknown. Among the traditionally conceived progenitors of western thought, Socrates is famed for his method. In summary form the socratic method consists of systematically posing an array of questions to an interlocutor with the aim of unsettling their prior assumptions and by doing so come to a deeper understanding. By relentlessly questioning assumed premises and deferring final judgement, one is kept in a state of continuous and suspended analytical deferment in which one is forced to strengthen one's thinking rather than rest lazily in a state of prejudice unaware of one's faulty presumptions. The Socratic method is dissolutive rather than solutional. It negates rather than posits, driving at the destruction of falsehoods by arguing down all other competing hypotheses rather than pushing for a presumed truth by simply reiterating justifications for one's pretenses.

Hintikka's work on what he calls socratic epistemology, which he examines

systematically in his work *Socratic Epistemology: Explorations of knowledge-Seeking and Questioning* and which this chapter is a response to, is an important contribution to the theory of inquiry. It bears instructive similarities with the present work. Hintikka also claims that epistemology as conventionally understood is flawed and partial. Its main concern is with knowledge as it is already possessed rather than as it is acquired, or in my phrasing, as it relates to ignorance. This parallels my position that epistemology is incomplete, although it falls short of the ultimate target. Where I disagree with Hintikka is that pinpointing knowledge (or information) acquisition as the missing piece of epistemology only goes half way. The real yin to the epistemological yang is the theory of ignorance, and the theory of inquiry is the middle term that ties them all together.

We also agree that reframing knowledge acquisition in terms of information acquisition is a more clearer concept to work with. However it is unclear to me if he supports replacing the concept of knowledge with information and can be spoken of interchangeably, an idea to which I am opposed because although the are related and knowledge can be said to be in some sense built out of information they are not the same. Also similar is his recommendation that the customary definitional goals of epistemology ("what is knowledge?") be sidelined in favor of an analysis of its application. knowledge is performed, and its ultimate test is how it influences decision making and achieves goals. For this reason Hintikka is also notable for his work on epistemological game theory. Furthermore, he too emphasizes the modal aspect of knowledge, but without the emphasis on entropy which I view as the indispensable grounding for the concept. knowledge permits one to justify disregarding possibilities inconsistent with it. It carves out the actual and the potential from the merely possible or the definitely impossible. This is in line with the negative idea of knowledge espoused here, that its function is to exclude irrelevant conclusions as much as it is to assert positivistic propositional truths, and that this is not merely an empty logical tautology.

These main features, the emphasis on the modal, game-theoretic, and informational aspects of epistemology, characterize Hintikka's approach.

Hintikka also offers an intriguing critique of Bayesian inference and the notion of cognitive bias which, while not central to his overall scheme, is apropos to how I use these concepts. He downplays their explanatory merit and asserts each to have serious conceptual faults. How these conclusions might entail not truly from accurate reasoning but falsely from ignorance blindness is the basis for a compelling counterargument.

Now to address Hintikka's views proper. Hintikka's analysis is modal in that the possession of knowledge dichotomizes world states into categories that designate admitted and excluded possible worlds. If we know there is a hurricane storming outside, we also know a chain of ampliative inferences that are connected to the possible world in which a storm is present. It would not be a good time to host a parade or having a picnic or fly a kite. The worlds in which those events might occur are not associated with a world in which a hurricane is happening. Certainly we can know that a flag on the flagpole outside is flapping and tree branches swaying in such a world.

While most of these inferences in this hurricane example are not strictly impossible, which might be construed to be the highest standard of modal exclusionary/inclusionary dichotomies, from a pragmatic standpoint they are not live options. It's not that we know it is a strictly impossible to fly a kite in a hurricane in any possible world, a standard of truth to which Lewis would hold absolute knowledge, knowledge of a consequence of natural law. A hurricane on a planet with different atmospheric properties may allow for a kite to fly. One can also tie a sturdy, reinforced kite to a steel cable attached to a car perhaps. In general though kite flying in a hurricane is not an action that would yield a typically desirable result. In that regard it certainly still is knowledge to affirm that a hurricane rules out kite flying.

In any case logical impossibility, as conventionally assumed, is inherently tautological, and as Hintikka argues, therefore informationally empty. So logical impossibility is of no consequence when making epistemic claims about empirical events and their significance.

Concerning Hintikka's work on epistemic game theory and epistemic logic I have no major comment or disagreement. I fundamentally agree that knowledge

is consequentialist in nature and best judged in terms of how it aids decisioning and boosts strategic reasoning about the practical effectiveness of inquiry. Game theory is a useful formalism for modeling investigative strategies. As for his formal epistemic logic, it is more of a tool for precisely designating epistemic propositions rather than an explanation for them. While this formal system may indeed be a powerful analytical device, it is not not an essential element to the theory of inquiry as it does not explain what inquiry is but rather clarifies how to think and speak about it.

The other main feature in common is the informational turn. Information is a useful conceptual shift for several reasons. Because information has a precise definition even if admittedly its common usage is often vague, a great deal of clarity is achieved. Extensive theoretical and mathematical understanding of information from the sciences can be recruited into epistemology. Emphasis on information, often sorely neglected in philosophy as Hintikka points out, even though it is all the rage in many other fields, also gives an excuse to bypass the controversial definitional problem of knowledge. It permits grounding vaguely defined knowledge on well-defined information. A unit of information is a bit, an utterly simple logical atom: yes or no, on or off. In chapter 7 of the book, Hintikka focuses more on the varied usages of information across multiple fields, points out the failure of philosophy to appreciate information as a useful concept, and recognizes that information cannot simply be smuggled into epistemology without explicitly laying out the justifications for doing so. Regardless, he concludes that "information, rather than knowledge or belief, should be the basic concept of epistemology." (p 190). The concept of information embeds a solid logical kernel into epistemology from which it can build out its formulations. It allows for a level of precision that and evades the eternal definitional problems of epistemology which have plagued it since its inception.

While I concur entirely that both modal and informational concepts are important, where I differ is how these concepts should be integrated. I believe these concepts say more about how the mind handles ignorance than it does playing the role of structuring knowledge. There is no contradiction here or sharp departure but more of a perspective shift that unveils a new way to think

about old problems. On my interpretation the modal aspect of knowing maps representations to information theoretic entropy which in turn correlates with physical entropy. This measure is an explicit proxy for the amount of uncertainty clouding a scenario. To the extent that knowledge dichotomizes possibilities into live or dead options, it parses excess entropic disorder as other-worldly possibilities of which we can safely remain ignorant and do not have to think about, leaving open those sets of live possibilities we describe by knowledge and infer by ampliative inference. On this point Hintikka and I are aligned, stating point blank, that information is connected with "classes of admitted and excluded possibilities. The more alternative possibilities a proposition excludes, the more narrowly we can restrict our attention, which is a manifestation of having more information" (p 191). What he falls short of realizing is how this is inherently bound to a theory of ignorance and overlooks the fundamental connection to entropy this modal perspective bears despite having his finger exactly on it.

Much in his analysis is spot on. What is curiously missing is any references to uncertainty, to information entropy, which is the unifying concept. Investigation is the reduction of information entropy and the balancing of an information gradient which results in a collapse of the range of plausible future states. In the same breath however, Hintikka rebukes Quine's misguided takes on the nature of information, which happen to be unaware of the central role of info-entropy in epistemology as well. While also agreeing that information is related to the sorting out of possibilities, Quine goes off the rails and asserts that for the notion of information to to be well-founded we would have to specify all possible states of the physical universe, a blatant impossibility which rules out the obvious practical applications of information theory that have been devised. Hintikka counters this by mentioning that, "a physical theory is typically applied, not to the entire huge universe as a whole, but to what physicists call "a system"— that is, a part (real or imaginary) sufficiently isolated to be considered a model of the theory in its own right." (p 191) The trouble here is with the phrase "sufficiently isolated." Quine and Hintikka take two polarized but similarly flawed views, either the entire universe must be counted the system, an impossible proposition,

or a manageable isolated system must. The latter is a system not so simple as to be trivial but not so complex it cannot be understood . But what exactly facilitates this "taking-as" that isolates the isolated system? Even taking the universe as the system is an act of isolating-isolating the universe as a whole apart from any reference to its parts as compartmentalized subsystems. What does it mean to "isolate" a system? Is anything in a physical world in which continuous fundamental fields pervade all matter truly isolable? At a certain point a decision must be made to draw a line and *to ignore* the rest of the universe beyond that boundary, even if it is pulling and tugging in minute or even substantial ways at the system more or less arbitrarily isolated by this commitment to a more narrowed object of attention. Thus to isolate is to implement structural ignorance. The act of isolating a system does not mean the system is in fact objectively isolated, there are in fact not definite empirical boundaries to any phenomenon. It is only described as bounded because otherwise it would be too difficult to think about. It is the mind that isolates.

I happen to side with Hintikka against Quine, for the sole reason that to take the entire universe as the system, is to take the system with necessarily the maximum entropy and therefore the highest degree of uncertainty (lowest information). Localized systems can in fact at least temporarily tend toward increasing order in contrast to its thermally disordered environment: we are living proof of that. Life in general is living proof of it. That is because technically speaking the second law applies only to closed systems, and the only system that is truly closed is the universe in its entirety. That is, it's closed under the first law, the law of conservation of energy whereas local region of the universe can get new energy from elsewhere as old energy dissipates from it. Any subsystem of the universe is in principle open, in the sense that energy flows can flow through it at a gradient. And it is precisely a difference in gradient which signals information. If the measure is the entire physical universe, it simply gives too much to sort through. The first step must be to truncate the universe through a brute act of simplification so that system under examination is intelligible.

This is all well and good, except that Hintakka's otherwise insightful analysis is lacking all references to uncertainty and entropy, which are fundamental to

information theory and therefore implicit to an informationally instructed epistemology. A type of ignorance structures knowledge, rather than simply and intuitively serving as its foil. To know is to collapse uncertainty and to recast possibilities into a narrowed range of probabilities. Inversely it is to exclude those determinations which are false or carry minimal information, to conclude, "I don't have to waste my time thinking about that. I can ignore that at no cost. I don't have to know about what doesn't matter." This is a distinct set of conclusions apart from, "I must take time to consider this, it could be important" because, to come to the determination that one must think about a proposition, an evaluative function must be applied to it. We do not cycle through every conceivable state of the world and evaluate each one to decide which is worth knowing about. We cut to the chase, cutting all ties off and closing the connection to the nonsensical and the remote. You can exclude a dead possibility without having to evaluate it but to think something through takes time and energy.

A bit of a digression is warranted here. Like many epistemic logicians, Hintikka embraces modal logic as a framework. Talk of possible worlds is a useful framing device to situate epistemic assertions. What exactly goes into this reasoning?

In a universe absolutely surfeited with variety and forever mixed up with itself, the most important aspect of any knowledge formation is distinctness. It's indeed the isolation of a system. From that act of discretizing the object of inquiry, its definite concept becomes thinkable. Yet this isolation cannot be derived from an act of pure knowing alone but also needs selective ignoring. My argument is that information is embodied in the decoupling of that isolated system from all the remaining totality of what is and what ifs, a detachment which is not inherently well defined, but which nevertheless draws a perimeter around the subject of analysis even if not a perfect circle.

Moving on to the issue of modal logic. Where does the idea of possibility come from? Loosely line with Lewis's conception of the principle of recombination, I suggest that the conception of possibility is derived from the observation of a world in which there is a vast entropic combinatorial space afforded by thermodynamic free variables, in which substitutions of elementary

states can be permutated arbitrarily. A universe that consists of two items \[A, B] and two locations \[1, 2] for each items, has one free variable in that the only possibility is for the items to remain where they are or to swap places. Therefore the state space of this simple universe is:

:: { [[A, 1] , [B,2]] | [[A,2] , [B,1]] } ::

With this combinatorial rule the roots of the concept of possibility can be traced to the abstraction of the sensory experience of the world of objects as an exponentially vast configuration space of substitutable items. However one defines an "item" is arbitrary, whether they are some kind of logical atom or perceptual objects, so long as there are enough of them to afford many combinations. Arguably the empirical source of the idea of possibility derives from generalizing the kinetic observation that some objects can be moved from one spot and replaced by another, or removed from a spot and put back in place, and this can be repeated with any number of objects. Once that observation has been made, the confirmation of a possibility is retroactively inferred. The observation can then be abstracted and applied to mental objects rather than the merely physical. The comparison of a memory of a world-state with a change in state marks out an actualized possibility. Such an event can happen because it did happen. Therefore, a class of events sharing the properties of that object must also have similar potential for recombination. To the extent that the mind can reason about the world accurately, it must then have the power to simulate this kinetic combinatorial behavior through counterfactual and hypothetical thinking. Once we have experienced the relational structure of the world, we can model it out in our minds without having to actually manipulate the things we are thinking about.

On the lowest level, the fact that objects can move and be rearranged is described by the statistical nature of heat physics, which assumes that everything is jostling next to each other and exchanging kinetic energy. It is precisely the displacement of objects caused by the spillover of potential energy that, when compared to a prior memory representation that is contrasted with it, signals information via surprisal. Which is to say the more things stay the same without rearranging, the lower the information the system has and the the more

predictable it becomes. How things differ at one point in time from how they were before captures information.

Having now reviewed and elaborated on the constituent elements of Hintikka's approach, what remains is to see how they are employed in his analysis and where it joins or departs from mine. My main criticism is that his method is rangebound and foreshortened by the "ignorance blindness" that has long afflicted epistemologists. He does not see ignorance and uncertainty as an epistemic force, commingling and adjacent to knowledge and forever dragging it along. Indeed the keywords appear preciously few times in the book and do not even show up in the index. This is a sign of the prejudice that ignorance is nothing, that it is purely the negation of knowledge. So what good does it serve thinking about it if it is just the negation of the "real thing"? All logics privilege the positive form as canonical, P comes before \~P. Thus ignorance is either ironically ignored and not mentioned at all, or is treated as an unimportant or unanalyzable externality to the project of epistemology, or a mere consequence of applying negation to the real object. None of this is true. Though in any inquiry the resolution of one's ignorance is the obvious goal of knowledge acquisition, Hintikka takes this for granted or does not acknowledge it at all. At most, Hintikka's treatment of ignorance is shallowly tacit. He does mention *bracketing*, a term borrowed from phenomenology that denotes a method used in the question-asking process. Bracketing sets aside uncertain answers and defers judgement on them until the requisite information has been acquired. Once further information clears them, the uncertain answers are retroactively unbracketed. Here Hintikka almost touches on the "inquisitive stance" described in the previous chapter. The ghost of the problem of ignorance seems to trial him without ever quite being named. He raises Meno's paradox, which I have elsewhere alluded to as a cognitive version of the computational halting problem as a use case of bracketing. The paradox is, "if you have what you're looking for, inquiry is unnecessary. If you don't know what you're looking for, inquiry is impossible." How can an inquirer know that they have reached the ultimate conclusion of or a fully actualized answer to an inquiry? I have already gone to some lengths to address this so I will not belabor it more than necessary.

Hintikka escapes it by asserting that the information gained from the inquiry at each stage allows the inquirer to refine the process by gaining sharper insight into it. This however does not resolve the initialization problem: how to start from a standpoint of obliviousness? The answer to me is quite intuitive, to realize one's obliviousness, even if the content of this present absence cannot be labeled. The way out is to intentionally objectify the related quantity of ignorance. To the extent that one wants to know something it can begin. Hintikka however seems to construct an elaborate system for dodging around this obvious answer. The desideratum of the inquiry, the goal state, is not a piece of knowledge one couldn't recognize in the first place as an adequate answer without reference to surefire evidence. It is the search for whatever may satisfy one's awareness of ignorance.

Granted as the inquiry progresses, that initially vague goal state is indeed successively approximated by the build up of ampliative inferences which were not present from the start. One merely needs an intimation that answers may reside in a particular direction. But if inquiry is a directed process, its goal cannot be knowledge, it cannot be something prior ignorance makes unrecognizable. Hintikka takes the desideratum, what I have previously referred to as the object of inquiry, as the payoff, the adequate answer that factually satisfies every condition of the question. In the initial stages, however, it is the absence of information that is the proper desideratum, the motivating cause, in that one continues investigating until one this vacuum is no longer apprehensible. It is not motivated by any achievable positive external datum, but by the desire to eliminate the consciousness of an information deficit, and this is true of an inquirer whether they indeed arrive at the true final answer and exhaust all relevant information or if they merely content themselves halfway there or even accept and believe falsehoods; it is true whether they are right or wrong that their ignorance has been satiated. "That's all I need (or want) to know" is not the same as "I know everything there is to know" but both are justifiable endpoints of an investigation given that ignorance is the prevailing condition in any case.

The criteria of satisfaction are self-provided and determined by judgement. This obviates Meno's paradox, as the supposed final desideratum of inquiry as

positive knowledge does not have to be obtained for certain. All that must be confirmed is that the irritation that isn one's self-consciousness of ignorance has subsided. This is true to the extent that belief formation and the satisfaction of doubt, as Peirce claimed, is the main goal of cognition rather than truth seeking. A truth seeker therefore just has higher standards for the revocation of their doubt, they postpone judgement and press on with their investigation well beyond the point that their own preconceptions have been gratified. Indeed Hintikka goes on to acknowledge that an inquiry is not necessarily a truth-seeking process. Information can be false, and the evaluative judgement of results is an independent working part separate from the process of information acquisition which itself only exhibits an eliciting function.

After Collingwood (1940), Hintikka finds credence in the notion that all questions stem from presuppositions that limit the scope of what may be asked, although he does take issue with aspects of the idea. A question always presumes a connection between its terms and imposes a structure which limits the range of available answers to it. For example, "What caused the dinosaurs to go extinct" implies that the dinosaurs went extinct, that dinosaurs existed in the first place, and perhaps if you are really unsparing, the concept of causality. No answer to this question which affirms the continued existence of dinosaurs, or denies them having existed , or denies they went extinct is permissible. That every question must call for a presupposition however makes little sense to me. (For what it's worth Hintikka seems to think that yes or no questions are presuppositionless because they are tautological and therefore logically empty.) To my thinking, for a presupposition to be the basis of a question it must be disbelieved. Unless there is an ulterior motive involved, for example one wishes to test what someone else knows by asking them a question and cares not for the answer, rarely are questions asked about which one already believes to be known. To hold a presupposition is to already assume one has the answer and takes it for granted. Presuppositions may be the result of answers from previous questions, in which case one may have formed enough of a picture of the subject to guide further inquiry. At the initial stages however, it is precisely the absence of presuppositions which get things moving. I have already made it abundantly

clear about what I believe triggers an inquiry.

Hintikka recognizes the problem of initiating an inquiry, and refers to Collingwood's assertion that all chains of presupposition trace back to an ultimate presupposition or initial premise. To me the content of this presupposition is obvious. It's merely the recognition "I don't know *this*." The use of the indexical "this" is essential because it picks out something specific. This is the most general form of a presupposition, from which any variety of more granular forms can be built out and refined as an information source is probed. Of course while "I don't know this" taken in isolation is nothing to work with, the recognition of it can initiate a process of speculation as to what actions might remedy it.

All told, I am willing to compromise with Hintikka and conclude that each question has a dual natured pairing of presuppositions. The assumptive nature of a question, which assumes the relations it expresses and imposes this form on the range of answers, has positive content. Yet there is also a negative presupposition at play which comes first, which serves as the setup function for the entire inquiry and which imposes its form on the positive presumption that flows from it.

In a similar vein is Hintikka's intriguing skepticism about Bayesian inference. Hintikka is suspicious that belief drives inquiry, placing it instead, after Peirce, at the end. The reason for this is that Bayesian reasoning presupposes an a priori selection problem of priors. To properly contextualize the input, Bayesian methods beg the question of how the right prior probabilities were chosen. There is, "…no presuppositionless Bayesian inference. Such an inference relies on a prior probability distinction that inevitably embodies assumptions concerning the world." (p. 201). Furthermore, to the extent that presumptive priors anchor the directionality of inference, circumscribing what counts as evidence for it or against it, there is always the chance that a chain of reasoning might be consistent with itself but wrong overall. "In the long evidential run, one's favorite measure of information might very well turn out to be dead wrong." (ibid).

The glaring flaw here is Hintikka's neglect of the principle of maximum entropy and the isomorphic concept of the uninformed prior (AKA diffuse prior)

used in Bayesian models. The uninformed prior is not exactly completely uniformed, it instead assumes the minimum of information as the prior probability distribution. It is often associated with Laplace's principle of indifference, which states in the absence of any evidence weighing way way or another, all possible conclusions should be assigned an equal probability. This forms the most "objective" starting point for a prior distribution without evidence. Initializing a Bayesian inference engine with an uninformed prior is akin to "looking at it with an open mind", of assuming the least about the data set and thereby avoid a biased model which would overly constrain the model. You might assume, after Quine, that this entails modeling the entire universe as that is the system with the maximum entropy, but in practice, the universe is never providable as a viable training input, and if it were, we would likely encounter the same problems as Wolpert's refutation of the feasibility of Laplacian Demon-like universal inference engine, the model itself being part of the universe. In any case the universe is unobservable in its totality, so it's an idle consideration to speculate further. The least is assumed about a data set for which assumptions can eventually be inferred.

One can find similarities in the use of a diffuse, uninformed prior with the notion of the intentional objectification of ignorance. Any inquiry by necessity begins with the minimum amount of information. What is the one piece of information of which the inquirer can be absolutely certain? That they don't know *something*, and that there must therefore be a something there is to know. From this vague, maximally ignorant initial state, more definite approximations may be sharpened down. The fact that the agent recognizes they don't know an inner something, immediately picks out an outer something to search for. It resolves the selection problem and orients the model toward a dataset that "fits." From this zeroed, null state, a more shapely characterization of the problem space can be developed. Another way of putting it is that while there is indeed no perfect rule to select a prior for the distribution, in the absence of a good practical or empirical reason to pick one over another, the prior which assumes the least and which has the potential to capture the most possible information from observation, leaving nothing out because everything is initially

admitted, is a solid fallback. It's akin to the idiom, "cast a wide net." This holds by the principle that if a solution resides somewhere between a minimum and maximum value, the assumption of an initial value closer to the maximum is more likely to encompass the solution somewhere under it than if the priors were under-specified and approximate the minimum.

Hintikka's critiques of Bayesian inference naturally flow into his critiques of cognitive bias, which he views as a wrongheaded idea. Indeed, he doubts their existence outright, going so far as to call it a "fad." In *Socratic Epistemology* he dedicates the entirety of chapter 9 to a takedown of the concept. He takes particular aim at Tvarsky and Kahneman's influential and original work on cognitive bias. Specifically he targets the conjunction fallacy as a crucial conceptual hinge on which cognitive bias depends and with the implication that debunking it also casts doubt on the remainder of the supposed biases.

In simplest terms, the conjunction fallacy is when two conditions are assumed to be more likely taken together than either taken separately. This is fallacious because the rules of probabilistic inference hold that in the absence of evidence, the assumption of a single condition is more likely than the conjunction of two or more, because this combines the uncertainties of each condition which self-evidently yields a less likely result than the uncertainty of a single condition.

The logical form of the conjunction fallacy is given as follows. Let's assume that S is a general description of the background evidence about B. Without explicitly stating either, S weakly implies C and D about B. Therefore which of these is more likely to be true of B:

1. C is true
2. C ∧ D is true

The fallacy comes from the notion that because S weakly implies both C and D, their conjunction must therefore be more probable even the though combined probabilities of two events cannot be greater than their independent probabilities. The more probabilities you add together, the more watered down the likelihood of their conjunction, but the human mind often seems to think that because the description S of B each loosely fits each of these conditions, their conjunction must constitute the most likely assumption.

The most famous example of the conjunction fallacy is the so called Linda problem. Consider the following scenario:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Based on this information, judge the likelihood of the following statements:

A: Linda is a bank teller (T)

B: Linda is a bank teller and active in the feminist movement (T ∧ F)

Tversky and Kahnemen found that the vast majority of survey respondents judged option B ((T ∧ F)) to be more likely despite its seeming incongruence with the laws of probability. Furthermore, further experiments in which expert statisticians were given the same problem yielded similar results, which led the researchers to theorize that the mind had innate biases that skew the judgement of probability.

Hintikka's argument against the conjunction fallacy essentially rests on the claim that the apparent bias is attributable to the description S constituting only a single information source, and that if multiple information sources are provided, conjunction is valid. Essentially, it is the same reasoning behind interviewing multiple witnesses in a court case. Each testimony is its own probability distribution, and as it turns out taking the geometric mean of multiple combined probability distributions leads to a more accurate result (Clemen and Winkler, 1999)[https://faculty.fuqua.duke.edu/\~clemen/bio/Published%20Papers/28.CombiningDistributions-Clemen&Winkler-RA-99.pdf]. Hintikka summarizes, "In a nutshell, I might sum up my diagnosis of the conjunctive fallacy by saying that given suitable collateral evidence, certain messages are self-confirming in that they tend to increase our confidence in their source (p 217). Supposing that answers to A and B, T and (T ∧ F) come from two different sources. As Hintikka explains, "The background information \[S] does not tell us whether T is likely to be true or not. If anything, the story about Linda makes it unlikely that she should be a bank teller… In contrast the information that F (…that Linda is active

in the feminist movement) is made likely by the background information… for this reason (T & F) coming from a different source of information might have to be assigned rationally a higher degree of credibility than the answer T, coming from a different witness." (p 216). That is to say, on basis of the background information, the most likely proposition taken independently is F, that Linda is active in the feminist movement, on the evidence that she was deeply concerned about issues of discrimination and social justice, and was politically active. This proposition however is not offered as a third choice. Option (T ∧ F), taken in conjunction with (T) from another source has more credibility than (T) taken by itself. This is a result of the fact that the evaluation of a probability is enhanced by multiple information sources.

This conclusion is a consequence of three requirements, Hintikka asserts:

1. The need for someone to modify their priors in light of new evidence
2. The fact that such a modification often amounts to a reevaluation of the degrees of reliability ascribed to different sources of information
3. The fact that the evidence conveyed by an answer may affect the very priors in relation to which the probability of the answer is to be judged

Given that the conjunction fallacy seem to disappear when multiple sources of information are included, it would appear that the fallacy is a byproduct of how the information is presented and the response options are provided within the experimental design, bias within the test rather than the subject. Were the conjunction fallacy instead an innate fault of reasoning, then the inclusion of multiple information sources oughtn't make out the choice of option (T ∧ F) to be rational. Were the conjunction fallacy real, we would expect it to apply in both cases. It should stick for both the minimum and maximum information cases, Of course, on its own Hintikka's explanation should be tested empirically to see how much it water it truly holds, but he is right in the sense that the underlying (un)reasoning involved in cognitive bias can be stripped out and analyzed purely in terms of the iron laws of logic from which people are claimed to deviate.

Implicitly, what Hintikka seems to be getting at here is that cognitive fallacies

disappear when we consider how the agent interacts with information sources in a plausible, ecologically valid way. Yet is it enough to go on? If some judgements can be inherently fallacious, Hintikka asks, "How come? How can it be that humans, including trained specialists, commit such elementary howlers? What is the psychological mechanism that creates them? In spite of the bulk of the relevant literature, I have not found a theoretically satisfying explanation of it." (p. 219) Indeed how come?

I have endeavored to find a preliminary answer to this question myself. In short form, there is no absolute necessity that human cognition as it its generally exhibited must exhibit perfect rationality. In fact from a biological perspective, such behavior ought to be more the exception than the rule, given that most of the bulk of the brain's architecture is a holdover from an ancestral, animalistic pre-rational time before the capacity for abstract thought emerged. And people, all of us, can and do display stunning irrationality. The question is if there are innate causes of these distortions. People can behave irrationally even if provided all the information and variety of information sources in the world. Humans were not evolved to be precise and meticulous reasoners. Least of all were they built to be intuitive probabilistic thinkers. They were designed to survive and reproduce, to win arguments in order to gain social status, and even to deliberately misrepresent circumstances to themselves if the resulting misbelief happens to benefit them. The psychological mechanism is evolutionary. It stems from the brain being a kluge, a loosely assembled collection of improvised subsystems rather than a finely tuned and honed instrument of ratiocination. If evolution happens to favor the irrational it will be prioritized in the thinking process over and above the rational. Hintikka seems to think that people in general hold themselves to the same standards as a logician such as himself. This is no the case. There are good reasons that structural ignorance is implicit to the operations of the brain and mind. There are tacit sufficient reasons behind the unreasonable, a purpose underlying how people can behave with seeming illogicality.

Ironically it deserves to be said that Hintikka himself falls for the fallacy of hasty generalization. On what grounds does he have to extend his -flawed-

critique of the conjunction fallacy to all cognitive biases in general? He is subject to his own lack of conditionalization of his priors and an absence of multiple information sources he himself is critiquing?

Ultimately, Hintikka reduces cognitive fallacies to the Bayesian prior probability selection problem, which as I have pointed out by taking the prior which assumes the least about the data, namely the uninformative prior. Given also the faulty assumption that multiple information sources, cross-checked against each other, disproves the existence of bias, and given also that "how come?" can be in fact be answered, the concept of cognitive bias arguably still stands.

Adding everything up, Hintikka's analysis of inquiry, while insightful, sophisticated and rigorous, is still confined to the epistemic positivism I'm criticizing as insufficient. In spite of his correctness that information acquisition and question-asking is an important component of a more complete epistemology, he only goes half way and his approach is thwarted by ignorance blindness. Absent in his analysis is an understanding of ignorance which is indispensable to understand the context in which inquiry unfolds.

This concludes what I'd consider to be the proper definition of the logical structure of inquiry, While I make no claims of definitiveness, I am confident that the features I have described are key to the appreciation of the inner mechanics of inquiry. While questioning is canonical facet of inquiry, it is by no means limited to just speech acts. Investigation is an embodied, extended process of information search and retrieval that involves everything that an agent can bring to bear in the quest to exhaust their surroundings of its secrets.

The Ignorance-first Nature of Inquiry

The Inquisitive Stance: Ignorance as Intentional Object

To understand inquiryvthe inquirer must first be understood. Inquiry unfolds with respect to an agent in relation to the unknown. Something special happens in the mind when we append a question mark to a sentence. The inquiring mood this signifies cannot be explained by an epistemology that lacks a proper

treatment of ignorance for the plain reason that the cause of the inquiry is not an object of knowledge but of ignorance. I call this mental process in which the object of ignorance is realized and pursued the *inquisitive stance* and identify it by several necessary attributes. In short, these defining features include *self-recognition of information inadequacy*, *the intentional objectification of ignorance*, *epistemic emotions (such as wonder, curiosity)*, *information seeking* and *information equilibrium*.

In broad terms, these different features of the inquiring process play specific roles at different stages. It must first begin with the self-recognition of an information deficit as this sets everything that follows in motion and one cannot ask about what one doesn't know one doesn't know. To pursue the which would serve as the replacement for this deficit, however, one must form some concept of it other than that it is an uncharacterizable brute unknown. One must establish "hooks" into it. One must mentally picture what one doesn't know in terms of what one doesn't, just as one can picture a missing jigsaw piece relative to all the pieces that have already been fitted. This hypothetical conceptualiziing is an intentional act in the technical sense that it is *about* something. To pursue an inquiry is to assume that whatever one doesn't know can be learned, that one's ignorance is tied to some as of yet hidden fact. Yet such simple ignorance does not resolve simply by realizing one has it. Inquiry must be motivationally driven to get one moving in search for information. This driver is curiosity and related epistemic emotions, which engage the agent to explore and probe their environment for clues. Once prompted by the desire to learn, next comes the proactive searching phase in which the inquirer seeks the appropriate information as befits their conception of their ignorance. Finally, as they draw learnings from this information source they reach a state in which the difference between external and internal information appears, their awareness of their ignorance evaporates, their curiosity dries up, or they are unable to find more information. Which is to say inquiry ceases at any point if one step in the chain is interrupted. At which time the inquirer no longer possess an intentional awareness of their ignorance of a given topic— and whether this perception is accurate is another case. Thus satisfied, they conclude their investigation either

believing everything they have learned is everything there is to learn or actually exhausting all unknowns relevant with the case sufficient to achieve their goal.

Taken together, these features comprise an engine of inquiry that functions to zero out a portion of ignorance that has been identified to a such a degree of salience that a question can be formed about it and appropriate information sources sought for it. This adoption of a conscious commitment to the belief that more is unknown than known and holding ones ignorance in mind as an intentional object, is the cognitive relation at the center of an epistemology of inquiry.

One last word of introduction. The inquisitive stance can be thought of as an explanation for the recursive interexchange between knowledge and ignorance mentioned in chapter 2 section 1. This basic relationship of exchange needs a mechanism of action. Ignorance does not transition into knowledge and knowledge back into ignorance by magic. There needs to be an underlying explanation. The key to this mystery is inquiry and its tools which are questions, that are like little hooks that latch onto the unknown and reel it in.

The process of inquiry does not begin with the question. The question is more a product of underlying mechanism. Indeed, the question is more toward the middle rather than the beginning of the process. The process begins with a metacognitive awareness of ignorance. To bother asking a question you must first realize there is something you don't know, and this something is definite enough that you can fit it into words. The awareness of this ignorance is what causes the question. The cause of inquiry, both its instigating factor and its final end goal, is thus a peculiar type of intentional object or "mental picture." The object of inquiry has no positive signification or presence initially, the kind of mental content that would come to mind if one were answering a question rather than asking it. Instead the intentional object which grounds an inquiry is a privative, an absence; a discontinuity in knowledge that is noticeable precisely for the fact that it is nothing, a gap that is demarcated by the edges of limited knowledge not lining up. The inquirer must fix this intentional object of ignorance in mind. This is what grounds the process of inquiry, initiating it, and providing the reference point from which it evolves. From there on forward, the inquirer monitors how

much this awareness of an absence diminishes as information is sought out that satisfies its incompleteness.

The saying "the absence of evidence is not evidence of absence" has a particular application here. Clearly some sort of negative evidence is required to come to the conclusion that one doesn't know something, but it is less clear what this evidence of absence consists of. I know with certainty that I cannot speak Chinese, but what exactly did I do to find out? It was an immediate and obvious conclusion, no kind of intensive review or thorough validation of my knowledge base was required to draw it. Presumably my evidence for knowing that I can't speak Chinese is found in the desire to speak it and then confirming that I am unable. Even then, my awareness of this ignorance is easier to grasp than that, as I make no attempt to speak it and do not need to try and fail to produce evidence that I can speak it to know I can't. Its absence from my knowledge is immediately apparent to me. Surely there are more ambiguous situations than confirming to myself that I don't know Chinese, lacunae within my knowledge that I am less conscious of and cannot even "demonstrate by failing to demonstrate" a lack of knowledge. Absence of knowledge is evidence of ignorance. The problem then becomes: how does one "gain purchase" into this ignorance, establish an adequate concept of it enough to ask a question that is a decisive step to overcome it?

Ignorance that can be questioned is of a more superficial kind than deep ignorance that is afflicted by *hypocognition*, a term from cognitive linguistics to describe the phenomenon of lacking words to describe a concept. Ignorance for which I can pose questions must be peripheral and mutable. It can be put in perspective, I can label its elements. Furthermore knowing that we don't know something by itself does not provide grounds for knowing what questions to ask about it. After all, we have nothing to work with. It is a raw, uninterpretable and absolutely simple datum: I don't know. Little can be done with that on its own. When we are truly stranded and stumped on an issue, we are stuck not only on the problem of our ignorance about the target domain, but also the secondary ignorance of not knowing which questions to ask that will pave a road out of the quagmire. If we are truly ignorant then hypocognition should come into play and

prevent any attempt to parse it into a directed question which points the way out. This poses a bit of a paradox, similar to Meno's. How can I know enough about what I don't know to ask the question that will lead to a correct answer and the resolution of my ignorance? This would seem to mean that I must at least be partially knowledgeable about the subject of questioning, otherwise I could not fit into words a good question to elicit the required information.

To be able to form a question suggests that the ignorance that stimulated it is somehow fragile or unstable, exposed to the possibility of knowledge, encompassed. Deep ignorance, ignorance of ignorance, the unknown unknown, is unquestionable. Not only is the content knowledge missing, but so is the judgement necessary to evaluate the meaning of this absence.

And it is precisely *judgement* which provides a way out of paradox. The question comes not from the ignorance of the content itself which is the first cause of it, and concerning which the question seeks an answer, but from the faculty of judgement that observes that ignorance. After all, nothing can come of nothing, a quantum of ignorance is dead weight and nothing can be done with it. Judgement, however, is a higher cognitive faculty which presides over and above the ignorance it evaluates. More must be said about how judgement relates to the intentionality of ignorance.

Ever since Brentano revived the concept from its medieval scholastic origins, a common assumption about intentionality is that it is "the hallmark of the mental" and that *all* mental states have this property of "aboutness." Furthermore it is assumed that this mental state refers to a content, which is to say, the object of intention cannot be null, it must have some "suchness" about it. "What is characteristic of every mental activity is, as I believe I have shown, the reference to something as an object." (Brentano, F., 1874 [1911, 1973], Psychology from an Empirical Standpoint, London: Routledge and Kegan Paul.) However, while maintaining that a mental state must have an object, Brentano also insists, and seemingly risks a contradiction by doing so, that this object need not necessarily exist outside the mind. "The terminus of the so-called \[intentional] relation does not need to exist in reality at all." This so called intentional inexistence was meant to highlight the argument that the mental object held its

own independent reality as an ontological fact. Whether or not it had a corresponding referent in the external world was optional. The idea of Bigfoot is its own fact, regardless of whether there is an external referent to which it attaches. Myths, dreams, fantasies, have some substantial reality even though they do not depend on external facts. Hence abstracts such as irrational numbers can still be entertained as factual entities without being traceable to an external referent that anchors its meaning.

All things said, ignorance exists. Despite it being a privative, a negative concept, this does not jeopardize its ontological status. Its existence is purely relational: it exists within the "interstices", the cracks between knowledge, in the same way the space between two buildings exists by virtue of the distance between them, but would not be said to exist without them. Awareness of ignorance is thus the awareness of the "presence of absence." It is the discovery that something is missing.

For judgement to take effect Brentano also held, it requires a presentation on which to operate: there must be a mental content to which judgement applies. Taken together these claims present an interesting and unusual case when interpreting ignorance as an intentional object.

Whether or not the intentional object has an external referent, it ought to be effable. It should be describable and in possession of properties that are enumerable. Even the paradigm of a privative concept, mathematical zero, has attributes. Zero isn't a pure nothing to which no more can be said. It is even, is neither positive nor negative, is an integer, is a multiple of all numbers, that when added or subtracted to a number always yields the same number, etc. Negative ignorance as an intentional object, however, has no such qualities. Certainly we can have thoughts that are about ignorance, yet these are of a special kind. This entire book is a demonstration of the fact that ignorance, fully conceived, has a detailed describable structure. It is not the always same for simple, negative ignorance, for simple not knowing, the self-identification of which is the first cause of inquiry. Surely this trouble does not show when one is positively ignorant, delusional or mistaken for example, as a mistaken or delusional belief is no different from a correct belief in terms of intentionality, it

has an object, wither existent or inexistent. A false belief is about something just as much as a true belief. The problem, precisely defined, is that when we have the belief that we are ignorant, there is nothing to which the object of this belief can be related in the mind, since we cannot represent to ourselves what we do not know.

The complication regarding intentional ignorance is directly related to Russell's work on establishing grounds for the justification of true negative existential beliefs. How can one assert a negative truth claim (e.g. Santa Claus does not exist), if there is no referent to represent what the proposition is about? If the referent designated by the proper name does not exist, then the name "Santa Claus" has no object, thus the proposition is not about anything at all. Reason holds that one cannot know a fact about the properties of something that doesn't exist, because what doesn't exist has no properties. Except the only means of determining facts is scanning the world for properties that obtain to it. If a proposition is about something, that object must exist. Russell escaped this riddle by holding that one must have *acquaintance* with- sensory experience of- a particular individual object to assign a logical proper name to it and therefore to meaningfully denote it (Russell, 1905). To denote the meaning of the name "Santa Claus", it must be recast as a description of a concrete individual. For example, "a jolly, bearded, red suit wearing man who lives in the North Pole and flies around in a magical sleigh to deliver presents to well-behaved children on Christmas." Thusly imbued with a definite description, the world of empirical objects can then be scanned to see if such an individual fitting that description exists or not. As such a definite description is then open to acquaintance while the proper name of a nonexistent object is not. It is impossible, so it is argued, to empirically establish the existence or nonexistence of what a proper name designates without a definite description of that object.

To tie this back to the problem of ignorance, no such definite description is possible for a single itemized unit or "quantum" of ignorance such as "I don't know Chinese." To provision a definite description of the object of ignorance is to know it by acquaintance, which is to thereby evaporate it as ignorance. Which is to say the more you describe what you don't know the more you know about it,

the more the object of ignorance to which your thought is fixed diminishes in proportion to which the knowledge about it grows. One cannot describe what you don't know in any detail without coming to know it.

The confusion over how one can picture one's ignorance well enough to find information that eliminates it subsides once self-recognized ignorance is understood as to be *metonymic*. To recognize one's ignorance is to substitute the empty privative (the absent which is present) for a kind of algebraic variable that stands for an as of unidentified or unformulated definite description of an object of knowledge. *It is to become aware of one's potential of knowledge*. It's like saying "*something* is going on outside", "*someone* will stop by to deliver the mail", or "do *whatever* it takes." These statements substitute an unidentified object with a generic term (something, someone, whatever) that functions as a kind of logical filler, expressing that there is in fact a referent that is predicated by the rest of the sentence even though it cannot be specified a proper name.

We do not have a clear idea of the intentional object of ignorance, because by definition we do not know the referent it connects to other than that this fact represents some unrealized potential knowledge. Yet the very recognition that one is ignorant of something proves that there must be something about which one is ignorant. To realize one's ignorance is to realize there is an order among all facts, and among the facts one knows, like jigsaw pieces, there must be at least one other. Otherwise, I would have to be acquainted with Chinese and able to render a description of it to know I don't know it, which is an absurdity. But self-recognized ignorance is always about *something*, there is some relation between the mental state that recognizes it and its corresponding unrealized object of knowledge. If one hears a loud bang outside one does not need to know what caused it for the hidden cause to be an intentional object that the thought is about. Clearly in this example such a hidden cause does not need to be effable for it to be the object of intention. By this means, judgement can "gain purchase" on a discrete "portion" of ignorance and hold it in mind as a distinct object. To raise ignorance into conscious awareness is to assign it a metonymic label which can then be fitted into propositions alongside positive knowledge.

All of this is admittedly esoteric, but it is necessary dirty work that must be sorted through to progress to a more well-founded explanation of the inquisitive process. All in all once ignorance has become the object of intention, deliberate, directed learning becomes a live option. One must know that it is a fact that one doesn't know something and that it is a fact that there is something one can know to pursue investigation. Without such intentionality, learning can only be accidental and passive, one cannot say to oneself, "I know what I must learn." Learning comes to you as a surprise rather than as the prize of a succesful hunt. The argument is that such goal driven learning cannot be engaged without first intentionally objectifying the ignorance connected to it. The hole must be found before it can be filled. Therefore to fix one's ignorance in mind is the first necessary step of proactive investigation.

On one reading, deliberate learning consists of two components that can be denoted by their respective inward and outward dispositions. The intrinsic component is this intentional objectification of ignorance described above. This discrete recognition of what one doesn't know results in its identification as a bounded finitude; it is to break off, in other words, a piece of quantified ignorance, carved at the joints, and detached from a vague and undifferentiated mass of deep ignorance about which there can be no such definite consciousness., to identify *an* unknown which can be directly targeted rather than *the* unknown which in its monolithic, featureless immensity is unassailable. Now, to mobilize learning, what helps it along is what has been called epistemic emotions such as curiosity or wonder. These emotions serve as motivators to compel the search for knowledge.

This brings us to the extrinsic component of deliberate learning. It is the process of active sensory exploration and searching for information sources within empirical space. In order to trigger the subjective conditions by which one judges one's ignorance to have been eliminated, some criteria of adequacy for external information uptake must be met. What are these criteria of satisfaction? The answer to the question must come from outside, and one must be reasonably confident that the answer one receives is sufficient. This raises Meno's paradox, how does one know one has learned enough? If, as we have assumed, the

awareness of one's ignorance is what compels directed inquiry, and if we assume that inquiry ceases once this awareness has dissipated because there is no more ignorance of which to be conscious, how can one know one has learned all there is to know without presuming that knowledge already? This sort of subjective halting problem has come up elsewhere, and the answer to it is not immediately obvious. There is no way to know before hand that the external information source is adequate, as this would presume knowledge it would be impossible to have unless one knew the answer already.

Even though one starts not knowing how to align the object of ignorance with an information source, the line must be drawn somewhere. You must start without necessarily knowing where to start. There are a few factors at play here: the depth and complexity of the sought for domain knowledge, the appropriateness of available information sources, and the precision with which the leading questions are defined. Now to address these in turn.

A complex subject obviously requires more cognitive effort to unpack. The overwhelmingness of its complexity may preclude asking the right questions at least until one has learned more about it through a kind of undirected and loose preliminary study of it. Its complexity may also increase the uncertainty that there is still more to find out, causing the opposite problem of when to end the investigation. A complex question is one that does not have a yes or no answer. There is in fact, uncertainty at both sides, at the start and at the end of the investigation and the appear to be related by the same problem. The only certainty one can have is that one's inner sense of ignorance, quantified and identified by one's questions, has run out because they've all been answered. Which is of course not an objective indication that everything there is to learn has been learned, only that the thing tying you to the information source has burnt up so that the connection has been severed and that your own subjective information needs have been met.

One solution to problem of determining the adequacy of information sources is that we rely on education and the authority of more informed people to guide learning. Using a curriculum, an arranged and prepared collection of information sources, cuts out the guesswork. Such shortcuts can evade the often aimless

"random walk" behavior of an unsystematized and unorganized study, a study that is lacking the right questions and therefore struggles to target the most valuable unknowns, the discovery of which would simplify and facilitate learning the rest. A simple example of this should suffice to demonstrate the principle. A child who wants to learn multiplication without first knowing they must learn addition will have a harder time, as multiplication is just the expedited form of addition of a number to itself as many times as the multiplier. For this reason addition is always taught before multiplication. In many domains of knowledge there are such "fundamentals" the learning of which accelerates learning the rest. Someone who is unable to mount a successful investigation is someone who is unable to identify which key unknowns they must direct their attention toward that unlock the potential to learn the rest.

Finally, a well-defined question clearly isolates the object of ignorance. Such that by comparing each term of the question and the answer proposition, it should be possible to tell by the form of the answer alone whether it satisfies the question. To use a trivial example one can intuitively know the question, "where is Mary and what is she doing?" is not satisfied by the answer "she is in the office." The response in this case obviously leaves the condition "and what was she doing?" dangling. After all the question is really a conjunction of two questions, "where was Mary" and "what is she doing" each of which must have an answer. Similarly, the response "anywhere but here," is inadequate. It fails to be disambiguate. Even "maybe she is in the office," fails to completely satisfy the requirements, because it expresses doubt on the part of the speaker. If an answer is to be accepted unconditionally, it must not add uncertainty to the questioner's uncertainty.

This last note about the formal symmetry between the terms of the question and the answer does not tell us about the accuracy of the semantic contents of the answer, but it does provide a basic analytic a priori check on the adequacy of the answer based on how their terms match up. Even "she is at the office" is a synthetic proposition that cannot be derived from the sentence alone and is only confirmable by experience. But the response that she is doing "anything" is a non-answer a priori, she is doing whatever one thing she is doing. A similar flaw

would appear if the response was " she is not at home and not in the office, but perhaps she is at the other places she usually goes." If the answer says too little or too much it's inadequate. The answer must fit the conditions of the question. If the answer does not reduce ambiguity but merely reproduces it, it is inadequate. Vagueness is unacceptable: *the answer must provide actionable information*. Like the notorious Glomar response made infamous by the CIA, it "neither confirms nor denies", failing to give actionable information by which the questioner can further narrow and target their information search. It's therefore self-evident that an undersupplied or oversupplied answer that does not align with the logical form of the question must be insufficient. On the contrary a modal question, for example Q: "what might Mary be doing" does afford a modal disjunctive response, A: "She might be doing x or y or z." Again, the point still holds that the question constrains the answer.

Similarly the complexity of the answer ought to be proportional to some degree to the question. This is somewhat analogous to Carl Sagan's famous quote, "extraordinary claims require extraordinary evidence." While it is certainly possible to resolve a detailed question with a simple yes or no, the probability of being able to do so diminishes as the complexity of the question and its answer criteria grows. This does not mean that a complex question is so because it is verbose, even a simple adverb can imply complexity: "is it raining out my window?" affords a simple yes or no answer, while "why is it raining…" demands a lesson in meteorology.

Yes or no questions are an important class of questions. They are recognized as a distinct type for good reason. What makes them special? It is that they are *purely logical* in that they exclusively accept the binary truth values of classical logic, true or false. They are the epitome of precisely defined. Such questions are strictly logically constrained to the fullest, forcing the answer to confine itself to an absolute commitment to one of two exclusive truth claims. Must every question reduce to simpler questions that accept a yes or no answer, as an information theorist might hope? A bit can be defined as a yes or no (digital) answer to an absolutely simple question, such as "is the lightbulb on". The game of 21 questions can be used to narrow down and identify an object from the

universe of all possible objects. This in itself is a paradigm example of a structured inquiry interrogating an information source and narrowing down the possibilities. I am unsure if in principle all answers can be reduced to a combination of yes or no atomic propositions, although this is certainly food for thought. As it turns out Hintikka (2007) suggests that perhaps every complex question can in fact be re-described in terms of yes or no questions, however this would entail an inefficient increase in verbiage as a general question is decomposed into a series each of which must be asked separately. In a class of 21 students, one day an apple on the teacher's table is stolen. The teacher can ask the whole class, "who stole the apple?" They could also ask each student individually "did you steal the apple (yes or no)?" While the precise binary nature of yes or no questions makes them appealing units for an imagined rigorous formal calculus of inquiry, it smacks to me more like wishful thinking that language would submit to such simplification without becoming distorted.

Moving on, to return to the problem of information adequacy, it is clarified further if we take a more humanized view of the investigation process and see how emotions play into it. The process is not a strictly "cold" logical and informational cognitive operation. As Helen de Cruz states, feelings such as awe and curiosity are epistemically productive. We do not require a formal definition of adequacy if we persist in being curious about a topic, in the words, if the investigation is no longer guided by definite questions or hypotheses but rather by a limitless desire to find out everything about a subject. Curiosity is learning for learning's sake. Wonder orients the agent toward open-ended discovery and gives a supra-rational reason to keep pushing for knowledge until the exhaustion of the emotional impulse. Emotions are not merely passive states that come over us, they are calls to action \([https://aeon.co/essays/how-awe-drives-scientists-to-make-a-leap-into-the-unknown]). Awe and curiosity permit the mind to open to new possibilities. Cruz calls this opening up of the mind cognitive accommodation. "The need for cognitive accommodation makes you aware that there is a lot you don't know…it is also an epistemic emotion, because it makes us aware of gaps in our knowledge." To consciously orient oneself to the unknown and to suppress the urge to fall back on preconceptions is to relax and

soften up hardened convictions and abandon illusions of knowledge completeness.

Psychologist Daniel Berlyne is noted for associating curiosity with a mood of deprivation and original work on the psychology of questioning. Uncertainty is often-but not always-unpleasant, and the recognition of it spurs a drive to eliminate it, which involves curiosity. Indeed several of his ideas anticipate facets of the intentional stance. He harkens back to the theories of Gestalt psychology and its "principle of closure" in which the mind "fills in the blanks" of an incomplete representation. Curiosity is uncontroversially what drives the impulse to close that gap. The idea in either case is that there is a gap, that it must be closed, and the question is how to go about it. In the 1954 paper *A Theory of Human Curiosity* Berlyne reviews the science of his day on the topic and settles on a conflict-driven theory of curiosity. Centering his theory on behaviorist analyses, questions are viewed as self-stimulating verbal behavior and the goal of curiosity is to remove an unpleasant stimulus. The centerpiece of his theory is that puzzling or unusual stimuli act as irritants the source of which is then actively explored.

(Footnote: Tangentially there are a few gripes I have with this work. Although the intentional object of ignorance is implicit in this "gap theory" the assumption is not that there is a gap within the representation, but nothing to start with. It is no filling in the blanks. The whole page is blank. The riddle is how novelty is overcome. Additionally, behaviorist conceptions of questioning, called thematic probes in the jargon, are considered to be the initiators of the exploratory behavior, which is self evidently absurd.)

([https://static1.squarespace.com/static/53a79084e4b01786c921de45/t/53a86486e4b009ec07711b59/1403544710847/A+Theory+of+Human+Curiosity+(Berlyne,+1954).pdf]). Berlyne also distinguished between epistemic curiosity—the urge to know ideas or abstractions— and perceptual curiosity, the desire to discern what is going on in the sensory environment. As it happens, more recent research suggests that epistemic curiosity its more associated with a positive sense of wonder while perceptual curiosity negative affect. ([https://www.psychologytoday.com/us/blog/why/201706/why-do-we-ask-why]) On

an animal level this makes some sense. The realm of abstract concepts may be confusing at times, but it can't hurt us, and it's hard for anybody not to find something that isn't immediately within range of the senses that worth's wondering about. Unknowns stirring in the physical environment however could be something dangerous or at least annoying.

Also intriguing is brain imaging studies which have associated curiosity with the nucleus accumbens, a structure strongly linked in reward and motivation. The nucleus accumbens is known to greatly involve dopamine in its mechanisms. It receives projections from dopamine-producing neurons, and the excitatory activations this induces is thought to play a significant role in experiences of reward, pleasure and motivation. ([https://pubmed.ncbi.nlm.nih.gov/25284006/])

The reason for mentioning this brain research is how the humble molecule dopamine and its role in the nucleus accumbens can explain so much. Dopamine is commonly referred to as the "pleasure molecule", implying that its expression is the biochemical source of pleasure, but this is misleading. Dopamine neurons not only respond when the animal receives a reward but also when a stimulus… predicts a reward. Such reward-predicting stimuli are conditioned rewards and have similar effects on learning and approach behavior to real rewards (Schultz, 2016). There is a sense then that dopamine is just as equally the curiosity molecule, or in the most general terms, the probability molecule. It has an essential role in orienting the brain toward stimuli which may be promising. In doing so it inherently encodes a probabilistic guess which carries an expectation of a possible reward. Precisely such an element would be required to orient the brain toward the pursuit of the unknown, especially if there is a payoff.

While the scientific research on curiosity and questioning is deep and interesting, my personal preference is that this is a subject better handled by philosophy. Specifically when it comes to questioning. Much is to be learned from the psychology of questions, but understanding them on the most fundamental level benefits from strictly viewing them as logical constructs.

These converging tendencies of self-recognized ignorance, the intentional objectification of ignorance and epistemic emotions, terminates in the decisive act

of posing a question. Philosopher Lani Watson has made a study out of questions. Watson is concerned with questioning questions. Such meta-questioning unveils a ripe space for thinking about thinking that has been surprisingly overlooked in the history of philosophy. If as is often said history repeats, or as the original Santayana quote reads, "those who cannot remember the past are doomed to repeat it," it can only mean that one has failed to step outside the first-order stream of events and neglects to adopt a second-order, distanced and removed perspective on it. Socrates, Watson points out, embodied the questioning spirit of philosophy through the Socratic method but declined to question this method of questioning itself. Fruitful recent work on the logic of questions, such as Hintikka (2007), tends to restrict itself to a logico-propositional analysis of questions as logical modality rather than as a more comprehensive cognitive psychological act. Watson goes beyond the propositional form of a question and reveals layers of subtlety in the act of questioning as a form of engagement with the world.

There are several findings to this end. A question, doesn't necessarily need a one-to-one answer. It may prompt many possibilities which, even in spite of their variety and in spite of no method for selecting the true answer from these options, still yields further information for further questioning. (Arguably this is not inconsistent with the previous discussion about the adequacy of answers, because the fact that an answer opens up many more possibilities for questioning means that it was inadequate.) What's more, Watson claims that a question does not require explicit linguistic expression and can be any information seeking intentional state and need not be verbalized. A google search query is an interesting case. One must simply have an unresolved inner state of uncertainty and seek to tap into an external information source that would reduce it. Of course verbal queries are more readily analyzable because of their formal, symbolic properties, but any information seeking behavior qualifies as a question broadly realized. By getting out of bed and opening up the window blinds to see if it is raining, one's behavior implicitly embodies an interrogative attitude without speech. As Watson asserts, "a question is an act." (Watson, Loni, What is a Question? The Philosophers Magazine, 82 (2018)). By extension, not every

answer has to be in a verbal form, it can be any type of confirmatory sense data. You do not need to rely on someone else to tell you what the weather is like, you can look outside.

In summary, all inquiry begins with the intentional recognition of one's ignorance as a metonymic psuedo-referent. A question or system of questions existentially quantifies or "draws a circle" so to speak around a portion of ignorance ("there is some x I do not know"), specifying it so that an actionable search for information sources can ensue. Epistemic emotions help spur this process on when the means for accurately evaluating the true state of one's ignorance—information incompletneess—relative to the complete objective information is lacking. These emotions are the experience of fascination: you keep learning without searching for answers. Investigation continues until the terms of the answer are judged adequate, and a subjective sense of parity or equilibrium of mutual information is achieved between the mind and the information source, that is, the information source is exhausted and the ignorance satisfied, like a cup of water being poured into another. Or it ends when the curiosity and drive for knowledge ceases. Inquiry proceeds to the extent that the decrease of ignorance is inversely proportional to the increase of questions answered until either one's awareness of one's ignorance or the information source dries up and no more can be found. Once one is no longer cognizant of one's ignorance as an intentional object, a feeling of satisfaction is induced and the inquiry ceases. This feeling of satisfaction is not necessarily accurate, however, because one may not be aware of the true extent of one's ignorance relative to what remains to learn. Deep ignorance-unknown unknowns —may still remain with respect to the topic after one has eliminated the intentional apprehension of surface ignorance,known unknowns.

This entire process can be illustrated by diagram:

(img)

Here is effectively the mapping of the codomain (image) of the function $f: I \rightarrow E$. Where each element of I is a state of ignorance that corresponds to an element of E which stands for an information state of an external information source. The dotted line surrounding both sets I and E represents the "horizon of

uncertainty" which envelops both domains. This is meant to represent that the cardinality of either set is undefined, and that true context of I or E may be greater than n. In other words there may be more external information En+y than what In selects as the initial baseline cutoff point. One's awareness of one's ignorance, may be less than the full extent. There may be more ignorance In+y than is encircled by the investigation. Thus, the sets are bordered by a fuzzy boundary. One cannot know the extent of one's ignorance and also therefore the extent of external information that is fit to answer to it. There must however be an initial state f(n) which initializes the function, a leading question as it were. The investigation begins within the horizon of uncertainty wherever a definite object of ignorance In can be identified. The double arrow stands for the bijective relation and feedback that occurs between both domains as n approaches 0. Like layers of an onion, the object of ignorance and the information source are peeled back until there is no more ignorance to be subdivided into more granular investigation states and no more external information to be unpacked. Each layer can be thought of a round of questioning and answering that successively iterates to the next step.

An answer to the question is simple if the question is simple. The question of whether it is raining outside can be confirmed by a simple yes or no answer by looking out the window. In this case, the depth of the inquiry function is f(0). Much more complex questions such as "what were the causes of World War One?" probably have a much greater depth.

This model of the relationship between the object of inquiry and an information source is of course idealized and simplified. However even where it fails to capture the true complexity of inquiry it can still be instructive. The spatial metaphor which views uncertainty as "atmospheric" and enveloping the domains is imperfect. Uncertainty does not simply exist on the outside, but also penetrates the process of learning in the course of its unfolding. One can imagine "ghost layers", in the information source, missing pages as it were, that halt or sidetrack the process. Imagine for instance a layer for Iy which has no corresponding layer Ey, and is instead replaced by another dotted line rectangle denoting a null layer of uncertainty. One may reach a given state in the function,

where the state of I has no corresponding state of E. There is no reason to assume that the process must be linear, no reason to assume every question must have an answer and one and only one answer, or that an information source is perfectly adequate to accomodate the inquiry. An investigation doesn't necessarily stop if in the course of the iterations of this cycle it encounters missing external information. Perhaps the solution is a recursive subroutine ("Why is that information missing?") which occurs in such a case. The missing information spurs a sub-investigation and creates a new object of inquiry, and the process is executed there until the missing information is returned before continuing to the higher level inquiry. In all the main idea is that there are coordinated relations between both the questions asked and the answers received that progress in step until the object of ignorance is completely eliminated while in proportion the information source is exhausted.

Inquiry is also readily diagrammable as a more familiar flowchart (expert system). A flowchart better illustrates the the conditional (if/then) logic of question and answering rather than a proportional extractive relationship between an information source and an inquisitive agent under ignorance. Say for example someone wants to know whether it is raining or snowing outside. Schematically they would proceed in this way:

Is it cloudy outside? (Y/N):

If N then : end query.

If Y then: is it warmer than 32 degrees?

———- If Y then: is it raining?

—- If Y then: end query

—If N then : is it snowing?

and so on…

Obviously, this toy example is not representative of how a person would actually approach checking the weather. They would skip to the chase and ask if it is raining or snowing, bypassing the assumed prerequisite of "is it cloudy." Or they would abstract the question of whether it was raining or snowing by asking what the weather is in general, which would give them the maximum of information with the most efficiency because they could also learn as a free

bonus whether it was sunny or just cloudy without precipitation all in one question. Or they'll just look outside and see for themselves. In reality not all questions are true or false, which further complicates matters. The fact that answers can have deeper semantics than yes or no vastly expands the amount of information that can be both expressed by and elicited by a given question. Even so, the logic of a flowchart captures the relational, treelike conditional reasoning of questioning.

For more complex investigations this formal structure of conditional dependencies is more relevant. A police detective engaged in interviewing a witness has no such heuristics. They cannot assume anything. They must start from the most general of questions and continue to specifics, trimming off branches of conditional possibility until they parse the one truth.

Here it gets more complicated. The conditional logic of questioning is not such a simple, stochastic process by which the next question to ask is dependent only on the current state and the previous answer. What must also be taken into consideration is the summation of prior answers everywhere in the tree. The simplified model assumes a perfect information source, which always gives the right answer at every state. If the information source has a probability of being wrong, however, the asker must factor in the posterior probabilities of their belief based on how well the answers to other questions are likely to infer it. (And of course it is a bit of a fantasy to suppose there are any information sources without a probability of being wrong.) Otherwise, they may end up asking the wrong questions based on wrong answers, which is a sure sign that an investigation has been derailed and is headed toward failure and the murky depths of positive ignorance. The goal of an investigation is truth, but the result may just as well end in falsehood. The tendency to view these phenomena in simplified, logically abstracted, skeletonized terms leads to idealizations that tend to ignore gritty practical realities. The possibility of being wrong or misled during investigation must be appreciated.

To validate any given answer for consistency, the truth values of the whole chain of questions everywhere in the tree must be summed together in a fashion reminiscent of Bayesian belief models in which the weights of various belief

states are readjusted depending on how much the actual input matches the expected value. In other words, one answer may fit the local requirements of its particular pathway down the tree, but it may conflict with an answer on an entirely different branch, so all data must be correlated. Granted, there are critics of this Bayesian modeling of question-asking, notably Popper and Hintikka, but as we will see Hintikka's critiques in specific are not strong enough to warrant discarding Bayesian inference as a useful model of probability expectations that is generalizable to inquisitive behavior.

To illustrate the point, consider someone who is caught in a lie. A liar must engage in a bit of double entry bookkeeping, monitoring both the reality and the fiction. To fish out the lie, an interrogator can ask questions that trap the liar into a logical corner of self-contradiction. A liar can keep up the facade so long as nothing else that they talk about, even seemingly tangential details, conflicts with the version of events to which they attest (or the all important evidence, which we will factor in soon). A crime suspect trying to save their own skin might lie in response to a given series of questions that has a specific assumption in mind. They may succeed to evade detection down this particular question path. Yet another chain of questions, indirectly related to the previous, can expose an inconsistency on the overall picture. The most explicit example of this dynamic is to separately interrogate both a suspect and a witness and to see how their narratives align. The witness, who is presumably neutral, can serve as a reference while the suspect's words are matched against theirs. Indeed such techniques for cornering a deceptive suspect into a contradiction are standard practice for detectives:

> Some experienced criminals or persons who have committed well-planned crimes believe that they can offer an alternate explanation for their involvement in the criminal event that will exonerate them as a suspect. An investigator may draw answers from this type of suspect by offering the same proposition that is offered for exoneration. This is the opportunity for a suspect to offer an alibi or a denial of the crime and an alternate explanation or exonerating evidence. It can be very difficult for a suspect to properly explain away all the evidence. Looking at the progression of the event, an interrogator can sometimes ask for additional

details that the suspect cannot explain. The truth is easier to tell because it happened, and the facts will line up. In contrast, a lie frequently requires additional lies to support the untrue statement. Examining a statement that is believed to be untrue, an interrogator can sometimes ask questions that expose the lies behind the original lie.

(Gehl, Rod & Plecas, Darryl. (2016). Introduction to Criminal Investigation: Processes, Practices and Thinking, chapter 8. New Westminster, BC: Justice Institute of British Columbia.)

The discussion up until now assumes that the interrogation process is entirely self-contained within the inner logic of the linguistic conversation. In practice, evidence is the decisive factor. Evidence can be thought of as what defines the grounds for belief of the investigator's question system that assign it satisfaction criteria. A shrewd investigator will base their questions on the evidence and they will judge the answers on the basis of the evidence, in conjunction with the answer system's logical self-consistency. These interactions between evidence and the verbal exchange together constitute an expanded framework of inquiry and are by no means independent. All told, the story the suspect tells must at first not contradict itself before it contradicts the evidence. Just like an alibi, evidence can be incomplete, fabricated, or misleading, and it is possible for a good liar to work around it. However, if they hoist themselves by their own petard it undermines the integrity of their rendition of events just as impactfully. And if they can't keep their story straight then how what they say compares to the evidence isn't even required to cast doubt.

The flowchart illustration sketched above is essentially the same as the concept of forward chaining, or in its more modern iteration, depth first search, used in some artificial intelligence models (Russell & Norvig 2009, p. 337). Originally, forward chaining was used in early AIs called expert systems. These programs encoded if-then representations between entities and were then fed input, which the system checked against its knowledge base. A medical expert system for example could be fed a list of symptoms and look up a diagnosis

match in its knowledge base. Researchers found that these implementations required excessive amounts of hardcoded, preinstalled data about relational entities, which effectively meant that their models were static, unscalable, and had no mechanism for learning. The world is constantly changing and new facts are discovered and old ones reevaluated. It would be impractical to continually manually update the expert system so that its knowledge base grew in step.

This fault of static representations in expert systems is related to the problem of meta-ignorance. From within a system of knowledge, how is it possible to judge that it is comprehensive? Such knowledge will only permit you to understand what it allows you to understand, enforcing a blindness on all else. How is it possible to know that a system of questions is sufficient or has been answered completely? Or alternatively, that a knowledge base is adequate to solve a novel problem? It is the riddle of epistemic incompleteness. As a Japanese proverb goes: a frog in the well knows not the ocean. How to know there isn't ever one more question to be asked, that no stones have been left unturned? Certainly one's background knowledge place limits on what questions can be broached. At a certain point a wall might be hit, at which time it may be not that everything all facts have been exhausted, but that one is mistaking the wall for the edge of the world. And when we are in this bind, it's hard to know which of these is the case because the information available, or unavailable as the case may be, applies equally to both. As Schopenhauer said, "every man takes the limits of his own field of vision for the limits of the world. This is an error of the intellect as inevitable as that error of the eye which lets us fancy that on the horizon heaven and earth meet." Such self-limitation means that while an inquiry will draw a line and stop, where that line stops might not coincide with the limit of what's knowable. Even the full mobilization of the powers of inquiry to conquer ignorance is ultimately swamped by it. In the long run entropy always prevails. How to avoid being the frog in the well, oblivious to the ocean?

The missing piece here and the key to defend against mistaking the well for the ocean is to expect that there is always at least one more stone left unturned. If there is indeed one, you will be more prepared for it, and if there isn't, there is no cost to this falsehood because you already knew everything anyway. Leave every

situation suspecting that you overlooked something. Of course one may be blind to one's own ignorance, and this pure ignorance, without any connection to knowledge to surface it from the abyss, may never come awareness. One may falsely believe an inquiry to be settled once curiosity has been satisfied or all questions one can think of are exhausted. It's better to always bet that there is still probably something you don't know even if not for any other reason but for the sheer mathematical fact that more is unknown than known in general. This fact can be used as a statistical base rate to which one can always refer to as an invariable guiding principle. (Indeed as was demonstrated in the previous section on computational limits, it is provable that the quantity of ignorance always exceeds that of knowledge.) It is revealing that the way forward for AI beyond expert systems was through the development of machine *learning*, as it is abundantly clear that learning is the ultimate adaptation to ignorance. A system starting off with minimal knowledge but a robust capacity to learn will eventually exceed a system, such as an expert system, which is preinstalled with a solid chunk of knowledge but no way to add to or revise it. This tacit incorporation of ignorance management into AI models has helped make them drastically more powerful in recent decades. In line with everything argued herein, these ML models suggest it's not what you know. It's what you can know, the openness to the universe of the unknown, which makes all the difference.

The solution to meta-ignorance is meta-cognition: the system of questions in turn can be subjected to questioning. One can ask oneself, "am I on the right track? Am I missing something?" As Peirce noted, beliefs are fixed— and thinking stops—when a certain level of doubt has been sated. On the flip side, this can be reversed, belief can dissolve back into doubt. In *The Fixation of Belief* Peirce takes a more unidirectional view of belief formation, assuming that doubt, through thinking, typically passes into belief but belief not back into doubt. People usually have a strong dictate for doubt and crave self-certainty. Yet doubt never loses its utility in prompting beliefs to be reevaluated and thus to enable learning. Such *second guessing* is not a sign of intellectual insecurity but a strategy to foolproof assumptions and catch errors.

When all these threads converge into one it becomes the inquisitive stance. In

summary, the inquisitive stance is the methodology by which a cognitive agent orients itself toward its own ignorance and renders it as a discrete object . Initiated by the intentional objectification of ignorance, driven by epistemic emotions and assisted by sensory exploration; directed by targeted questions and confirmed by answers, the agent seeks information which counterbalances this initial deficit. A kind of work cycle is in play, in which information flows across a gradient, from the data rich outside world into the ignorant mind. It's a deceptively powerful statement to say, "I don't know *that*." The living possibility, or at least the enhanced likelihood, of obtaining that knowledge is a direct consequence of this realization. Without consciousness of one's ignorance as a distinct object- although an empty, null object, the presence of absence- the only other way to discover what one doesn't know is by accident or caprice. Unless a deliberate project to banish recognized ignorance is undertaken, the diminishment of it is contingent on passive experience and incidental encounters.

Often the pursuit of knowledge is held up as one of the most noble aspirations of which humanity is capable. Yet this is not the right way to look at it. The knowledge is not "out there" to be pursued. At best, only information can be said to exist externally, but information is not knowledge. knowledge exists only in the mind, and it does not exist until the mind searching for it has derived it. So what exactly is being pursued? You can't find what cannot be found. The same goes for the term "knowledge acquisition", although for the sake of convention I still use it. What is there to acquire? The conclusion comes before the premise. This is circular.

It is more meaningful to speak of fleeing from ignorance. The goal is not to pursue knowledge but to escape being ignorant, and these are different concepts. By attending to ignorance we necessarily hone our conception of it, highlight its outlines, identify the shape of the missing puzzle piece from the arrangement of all the surrounding pieces. We find out what we need to move away from. Information may be the prerequisite of knowledge, but all the information in the world does make a difference unless the keyhole in which it fits is at first located. It's more accurate and insightful to say we flee ignorance than pursue

knowledge. It is by getting away from our ignorance, of which we have become so aggravatingly conscious, that we thrust ourselves toward information which is adequate to resolve it, and as this information is absorbed it is constructed into knowledge. As has hopefully been argued to satisfaction, this inversion is not a mere tautology, a way of saying the same thing in double-negated terms. Ignorance is the starting point, the predominating epistemic condition, and knowledge is a mere reaction to it. knowledge is never as certain as ignorance. If nothing else we can be certain of some form of ignorance. knowledge in contrast can be fragile, reversible; it can expire, degrade, and be superseded. Contingent, conditional knowledge may be falsified as soon as the conditions on which it depends change. Not so ignorance conceived in its totality. Indeed it is much easier to definitely define ignorance than it is knowledge, which may explain why the latter has remained an ongoing challenge. Perhaps this suggests something about knowledge's peripheral nature. Is knowledge so hard to define because it is inherently mutable, always shadowed, bordered, infused and led on by ignorance? What passed for knowledge in the now outdated science of the 1800s is no longer considered knowledge, but who are we to conclude that the present state of knowledge is on a more solid footing?

Meanwhile ignorance remains. While some of its territory has been conquered, it seems to grow an acre for every acre claimed. If we allow that knowledge as I believe we should that does not have to be perfect and immutable, that it can be provisional, logically fuzzy, and pockmarked by ignorance, then what we are simply doing is crediting ignorance its proper place in epistemology by subordinating knowledge to it. As it happens, this is also the route to intellectual humility, the understanding that the world is more complex than we think it is. To the extent that knowledge can be said to be possessed, it is in that it reorders and clarifies the knower's relationship with their own ignorance, allows them to adapt. But just as the size of any countable subsequence of an infinite uncountable series, no matter how large, is always infinitely less than the size of the uncountable series, so too is knowledge subsumed by ignorance.

Being asked questions is important to the theory of inquiry too. A question

which cannot be answered is also a question which cannot be understood. An answerable question assumes a parity between the known and the unknown, it assumes knowability. An unanswerable question, even if it is well formed, still yields some information in the same way that null results do in science. It still points to something, and raises the further question about why exactly this question can't be answered. It shows us dead ends and therefore helps reveal which paths remain viable. (Such dead ends can be "bracketed" as pending and uncertain and unbracketed later when more information comes to light, as they may only be unanswerable given the present state of inquiry.) By going unanswered a stumping question tells us what we can't know, it provides proof of inadequacy which again furthers investigation. It validates the integrity of the knowledge base. An unanswerable well-formed question raises the specter of structural ignorance, which is reflected not by simple lack of knowledge but by the limitations of the framework of existing knowledge. knowledge which prioritizes certain focuses and so neglects others and is inherently circumscribed. Newtonian physics is certainly knowledge, but it still falls short of the fuller picture, and incorporating relativistic principles into it —incompatible knowledge forms— unravels it.

For these reasons it is the questioners, who are by extension always learners rather than lazily self-satisfied knowers, who are the real epistemic powerhouses. They do not stop at an illusory boundary of falsely complete knowledge, the telltale sign of epistemic senescence, but continually push the envelope into the interminable unknown. They reside on the edge, continually revolving, turning over their minds, adapting their ideas to new discoveries. They anticipate complexity, and if they overshoot their estimation of it, they are still better off than if they underestimated it. They situate their thinking at the point of intersection between the edge of knowledge and the beginning of the unknown, at that one vital point where the two curves overlap. The best way to demystify these ancient philosophical enigmas about how we know is exactly this inquisitive stance in which the unknown, punctured by inquiry, continually flows into the known by having its information content siphoned off into it by the asking of penetrating questions.

Hintikka's Socratic Epistemology

Questions are doubtlessly an essential feature of the learning and thinking process. Interrogative sentences were used as far back as the Epic of Gilgamesh in 21BCE, and while it can't be proven it's a safe bet that for as long as there has been language there has been interrogative speech acts. In all likelihood, curiosity as the primal emotive precursor of verbal questioning predates humans, and perhaps even primates. Chimpanzees exhibit behavior that can be unambiguously interpreted as expressions of curiosity. Chimpanzees who have been exposed to the sight of fire for the first time have been observed to display investigative behavior such as fixating their attention on it and using tools like sticks to poke and prod it (Byrne, 2013). Ants, in their restless drive to scan the area round the colony for food sources may not be exhibiting conscious curiosity, but they are behaving that way for the same reasons that curiosity exists in primates.

As fundamental epistemic constructs, questions and the process of asking them has understandably been the primary stimulus behind the history of ideas. After all you can have no development of ideas unless there is something to trigger change and the incorporation of new information. And rather than to wait to be proven wrong one can actively question and take control of one's ignorance. Such changes, if not always preceded by a question, in every case in which it is always stems from a deliberate engagement with the unknown. Among the traditionally conceived progenitors of western thought, Socrates is famed for his method. In summary form the socratic method consists of systematically posing an array of questions to an interlocutor with the aim of unsettling their prior assumptions and by doing so come to a deeper understanding. By relentlessly questioning assumed premises and deferring final judgement, one is kept in a state of continuous and suspended analytical deferment in which one is forced to strengthen one's thinking rather than rest lazily in a state of prejudice unaware of one's faulty presumptions. The Socratic method is dissolutive rather than solutional. It negates rather than posits, driving at the destruction of falsehoods by arguing down all other competing hypotheses

rather than pushing for a presumed truth by simply reiterating justifications for one's pretenses.

Hintikka's work on what he calls socratic epistemology, which he examines systematically in his work *Socratic Epistemology: Explorations of knowledge-Seeking and Questioning* and which this chapter is a response to, is an important contribution to the theory of inquiry. It bears instructive similarities with the present work. Hintikka also claims that epistemology as conventionally understood is flawed and partial. Its main concern is with knowledge as it is already possessed rather than as it is acquired, or in my phrasing, as it relates to ignorance. This parallels my position that epistemology is incomplete, although it falls short of the ultimate target. Where I disagree with Hintikka is that pinpointing knowledge (or information) acquisition as the missing piece of epistemology only goes half way. The real yin to the epistemological yang is the theory of ignorance, and the theory of inquiry is the middle term that ties them all together.

We also agree that reframing knowledge acquisition in terms of information acquisition is a more clearer concept to work with. However it is unclear to me if he supports replacing the concept of knowledge with information and can be spoken of interchangeably, an idea to which I am opposed because although the are related and knowledge can be said to be in some sense built out of information they are not the same. Also similar is his recommendation that the customary definitional goals of epistemology ("what is knowledge?") be sidelined in favor of an analysis of its application. knowledge is performed, and its ultimate test is how it influences decision making and achieves goals. For this reason Hintikka is also notable for his work on epistemological game theory. Furthermore, he too emphasizes the modal aspect of knowledge, but without the emphasis on entropy which I view as the indispensable grounding for the concept. knowledge permits one to justify disregarding possibilities inconsistent with it. It carves out the actual and the potential from the merely possible or the definitely impossible. This is in line with the negative idea of knowledge espoused here, that its function is to exclude irrelevant conclusions as much as it is to assert positivistic propositional truths, and that this is not merely an empty

logical tautology.

These main features, the emphasis on the modal, game-theoretic, and informational aspects of epistemology, characterize Hintikka's approach.

Hintikka also offers an intriguing critique of Bayesian inference and the notion of cognitive bias which, while not central to his overall scheme, is apropos to how I use these concepts. He downplays their explanatory merit and asserts each to have serious conceptual faults. How these conclusions might entail not truly from accurate reasoning but falsely from ignorance blindness is the basis for a compelling counterargument.

Now to address Hintikka's views proper. Hintikka's analysis is modal in that the possession of knowledge dichotomizes world states into categories that designate admitted and excluded possible worlds. If we know there is a hurricane storming outside, we also know a chain of ampliative inferences that are connected to the possible world in which a storm is present. It would not be a good time to host a parade or having a picnic or fly a kite. The worlds in which those events might occur are not associated with a world in which a hurricane is happening. Certainly we can know that a flag on the flagpole outside is flapping and tree branches swaying in such a world.

While most of these inferences in this hurricane example are not strictly impossible, which might be construed to be the highest standard of modal exclusionary/inclusionary dichotomies, from a pragmatic standpoint they are not live options. It's not that we know it is a strictly impossible to fly a kite in a hurricane in any possible world, a standard of truth to which Lewis would hold absolute knowledge, knowledge of a consequence of natural law. A hurricane on a planet with different atmospheric properties may allow for a kite to fly. One can also tie a sturdy, reinforced kite to a steel cable attached to a car perhaps. In general though kite flying in a hurricane is not an action that would yield a typically desirable result. In that regard it certainly still is knowledge to affirm that a hurricane rules out kite flying.

In any case logical impossibility, as conventionally assumed, is inherently tautological, and as Hintikka argues, therefore informationally empty. So logical impossibility is of no consequence when making epistemic claims about

empirical events and their significance.

Concerning Hintikka's work on epistemic game theory and epistemic logic I have no major comment or disagreement. I fundamentally agree that knowledge is consequentialist in nature and best judged in terms of how it aids decisioning and boosts strategic reasoning about the practical effectiveness of inquiry. Game theory is a useful formalism for modeling investigative strategies. As for his formal epistemic logic, it is more of a tool for precisely designating epistemic propositions rather than an explanation for them. While this formal system may indeed be a powerful analytical device, it is not not an essential element to the theory of inquiry as it does not explain what inquiry is but rather clarifies how to think and speak about it.

The other main feature in common is the informational turn. Information is a useful conceptual shift for several reasons. Because information has a precise definition even if admittedly its common usage is often vague, a great deal of clarity is achieved. Extensive theoretical and mathematical understanding of information from the sciences can be recruited into epistemology. Emphasis on information, often sorely neglected in philosophy as Hintikka points out, even though it is all the rage in many other fields, also gives an excuse to bypass the controversial definitional problem of knowledge. It permits grounding vaguely defined knowledge on well-defined information. A unit of information is a bit, an utterly simple logical atom: yes or no, on or off. In chapter 7 of the book, Hintikka focuses more on the varied usages of information across multiple fields, points out the failure of philosophy to appreciate information as a useful concept, and recognizes that information cannot simply be smuggled into epistemology without explicitly laying out the justifications for doing so. Regardless, he concludes that "information, rather than knowledge or belief, should be the basic concept of epistemology." (p 190). The concept of information embeds a solid logical kernel into epistemology from which it can build out its formulations. It allows for a level of precision that and evades the eternal definitional problems of epistemology which have plagued it since its inception.

While I concur entirely that both modal and informational concepts are important, where I differ is how these concepts should be integrated. I believe

these concepts say more about how the mind handles ignorance than it does playing the role of structuring knowledge. There is no contradiction here or sharp departure but more of a perspective shift that unveils a new way to think about old problems. On my interpretation the modal aspect of knowing maps representations to information theoretic entropy which in turn correlates with physical entropy. This measure is an explicit proxy for the amount of uncertainty clouding a scenario. To the extent that knowledge dichotomizes possibilities into live or dead options, it parses excess entropic disorder as other-worldly possibilities of which we can safely remain ignorant and do not have to think about, leaving open those sets of live possibilities we describe by knowledge and infer by ampliative inference. On this point Hintikka and I are aligned, stating point blank, that information is connected with "classes of admitted and excluded possibilities. The more alternative possibilities a proposition excludes, the more narrowly we can restrict our attention, which is a manifestation of having more information" (p 191). What he falls short of realizing is how this is inherently bound to a theory of ignorance and overlooks the fundamental connection to entropy this modal perspective bears despite having his finger exactly on it.

Much in his analysis is spot on. What is curiously missing is any references to uncertainty, to information entropy, which is the unifying concept. Investigation is the reduction of information entropy and the balancing of an information gradient which results in a collapse of the range of plausible future states. In the same breath however, Hintikka rebukes Quine's misguided takes on the nature of information, which happen to be unaware of the central role of info-entropy in epistemology as well. While also agreeing that information is related to the sorting out of possibilities, Quine goes off the rails and asserts that for the notion of information to to be well-founded we would have to specify all possible states of the physical universe, a blatant impossibility which rules out the obvious practical applications of information theory that have been devised. Hintikka counters this by mentioning that, "a physical theory is typically applied, not to the entire huge universe as a whole, but to what physicists call "a system"— that is, a part (real or imaginary) sufficiently isolated to be considered a model of the

theory in its own right." (p 191) The trouble here is with the phrase "sufficiently isolated." Quine and Hintikka take two polarized but similarly flawed views, either the entire universe must be counted the system, an impossible proposition, or a manageable isolated system must. The latter is a system not so simple as to be trivial but not so complex it cannot be understood . But what exactly facilitates this "taking-as" that isolates the isolated system? Even taking the universe as the system is an act of isolating-isolating the universe as a whole apart from any reference to its parts as compartmentalized subsystems. What does it mean to "isolate" a system? Is anything in a physical world in which continuous fundamental fields pervade all matter truly isolable? At a certain point a decision must be made to draw a line and *to ignore* the rest of the universe beyond that boundary, even if it is pulling and tugging in minute or even substantial ways at the system more or less arbitrarily isolated by this commitment to a more narrowed object of attention. Thus to isolate is to implement structural ignorance. The act of isolating a system does not mean the system is in fact objectively isolated, there are in fact not definite empirical boundaries to any phenomenon. It is only described as bounded because otherwise it would be too difficult to think about. It is the mind that isolates.

I happen to side with Hintikka against Quine, for the sole reason that to take the entire universe as the system, is to take the system with necessarily the maximum entropy and therefore the highest degree of uncertainty (lowest information). Localized systems can in fact at least temporarily tend toward increasing order in contrast to its thermally disordered environment: we are living proof of that. Life in general is living proof of it. That is because technically speaking the second law applies only to closed systems, and the only system that is truly closed is the universe in its entirety. That is, it's closed under the first law, the law of conservation of energy whereas local region of the universe can get new energy from elsewhere as old energy dissipates from it. Any subsystem of the universe is in principle open, in the sense that energy flows can flow through it at a gradient. And it is precisely a difference in gradient which signals information. If the measure is the entire physical universe, it simply gives too much to sort through. The first step must be to truncate the universe through a

brute act of simplification so that system under examination is intelligible.

This is all well and good, except that Hintakka's otherwise insightful analysis is lacking all references to uncertainty and entropy, which are fundamental to information theory and therefore implicit to an informationally instructed epistemology. A type of ignorance structures knowledge, rather than simply and intuitively serving as its foil. To know is to collapse uncertainty and to recast possibilities into a narrowed range of probabilities. Inversely it is to exclude those determinations which are false or carry minimal information, to conclude, "I don't have to waste my time thinking about that. I can ignore that at no cost. I don't have to know about what doesn't matter." This is a distinct set of conclusions apart from, "I must take time to consider this, it could be important" because, to come to the determination that one must think about a proposition, an evaluative function must be applied to it. We do not cycle through every conceivable state of the world and evaluate each one to decide which is worth knowing about. We cut to the chase, cutting all ties off and closing the connection to the nonsensical and the remote. You can exclude a dead possibility without having to evaluate it but to think something through takes time and energy.

A bit of a digression is warranted here. Like many epistemic logicians, Hintikka embraces modal logic as a framework. Talk of possible worlds is a useful framing device to situate epistemic assertions. What exactly goes into this reasoning?

In a universe absolutely surfeited with variety and forever mixed up with itself, the most important aspect of any knowledge formation is distinctness. It's indeed the isolation of a system. From that act of discretizing the object of inquiry, its definite concept becomes thinkable. Yet this isolation cannot be derived from an act of pure knowing alone but also needs selective ignoring. My argument is that information is embodied in the decoupling of that isolated system from all the remaining totality of what is and what ifs, a detachment which is not inherently well defined, but which nevertheless draws a perimeter around the subject of analysis even if not a perfect circle.

Moving on to the issue of modal logic. Where does the idea of possibility come from? Loosely line with Lewis's conception of the principle of

recombination, I suggest that the conception of possibility is derived from the observation of a world in which there is a vast entropic combinatorial space afforded by thermodynamic free variables, in which substitutions of elementary states can be permutated arbitrarily. A universe that consists of two items \[A, B] and two locations \[1, 2] for each items, has one free variable in that the only possibility is for the items to remain where they are or to swap places. Therefore the state space of this simple universe is:

:: { [[A, 1] , [B,2]] | [[A,2] , [B,1]] } ::

With this combinatorial rule the roots of the concept of possibility can be traced to the abstraction of the sensory experience of the world of objects as an exponentially vast configuration space of substitutable items. However one defines an "item" is arbitrary, whether they are some kind of logical atom or perceptual objects, so long as there are enough of them to afford many combinations. Arguably the empirical source of the idea of possibility derives from generalizing the kinetic observation that some objects can be moved from one spot and replaced by another, or removed from a spot and put back in place, and this can be repeated with any number of objects. Once that observation has been made, the confirmation of a possibility is retroactively inferred. The observation can then be abstracted and applied to mental objects rather than the merely physical. The comparison of a memory of a world-state with a change in state marks out an actualized possibility. Such an event can happen because it did happen. Therefore, a class of events sharing the properties of that object must also have similar potential for recombination. To the extent that the mind can reason about the world accurately, it must then have the power to simulate this kinetic combinatorial behavior through counterfactual and hypothetical thinking. Once we have experienced the relational structure of the world, we can model it out in our minds without having to actually manipulate the things we are thinking about.

On the lowest level, the fact that objects can move and be rearranged is described by the statistical nature of heat physics, which assumes that everything is jostling next to each other and exchanging kinetic energy. It is precisely the displacement of objects caused by the spillover of potential energy that, when

compared to a prior memory representation that is contrasted with it, signals information via surprisal. Which is to say the more things stay the same without rearranging, the lower the information the system has and the the more predictable it becomes. How things differ at one point in time from how they were before captures information.

Having now reviewed and elaborated on the constituent elements of Hintikka's approach, what remains is to see how they are employed in his analysis and where it joins or departs from mine. My main criticism is that his method is rangebound and foreshortened by the "ignorance blindness" that has long afflicted epistemologists. He does not see ignorance and uncertainty as an epistemic force, commingling and adjacent to knowledge and forever dragging it along. Indeed the keywords appear preciously few times in the book and do not even show up in the index. This is a sign of the prejudice that ignorance is nothing, that it is purely the negation of knowledge. So what good does it serve thinking about it if it is just the negation of the "real thing"? All logics privilege the positive form as canonical, P comes before \~P. Thus ignorance is either ironically ignored and not mentioned at all, or is treated as an unimportant or unanalyzable externality to the project of epistemology, or a mere consequence of applying negation to the real object. None of this is true. Though in any inquiry the resolution of one's ignorance is the obvious goal of knowledge acquisition, Hintikka takes this for granted or does not acknowledge it at all. At most, Hintikka's treatment of ignorance is shallowly tacit. He does mention *bracketing*, a term borrowed from phenomenology that denotes a method used in the question-asking process. Bracketing sets aside uncertain answers and defers judgement on them until the requisite information has been acquired. Once further information clears them, the uncertain answers are retroactively unbracketed. Here Hintikka almost touches on the "inquisitive stance" described in the previous chapter. The ghost of the problem of ignorance seems to trial him without ever quite being named. He raises Meno's paradox, which I have elsewhere alluded to as a cognitive version of the computational halting problem as a use case of bracketing. The paradox is, "if you have what you're looking for, inquiry is unnecessary. If you don't know what you're looking for,

inquiry is impossible." How can an inquirer know that they have reached the ultimate conclusion of or a fully actualized answer to an inquiry? I have already gone to some lengths to address this so I will not belabor it more than necessary. Hintikka escapes it by asserting that the information gained from the inquiry at each stage allows the inquirer to refine the process by gaining sharper insight into it. This however does not resolve the initialization problem: how to start from a standpoint of obliviousness? The answer to me is quite intuitive, to realize one's obliviousness, even if the content of this present absence cannot be labeled. The way out is to intentionally objectify the related quantity of ignorance. To the extent that one wants to know something it can begin. Hintikka however seems to construct an elaborate system for dodging around this obvious answer. The desideratum of the inquiry, the goal state, is not a piece of knowledge one couldn't recognize in the first place as an adequate answer without reference to surefire evidence. It is the search for whatever may satisfy one's awareness of ignorance.

Granted as the inquiry progresses, that initially vague goal state is indeed successively approximated by the build up of ampliative inferences which were not present from the start. One merely needs an intimation that answers may reside in a particular direction. But if inquiry is a directed process, its goal cannot be knowledge, it cannot be something prior ignorance makes unrecognizable. Hintikka takes the desideratum, what I have previously referred to as the object of inquiry, as the payoff, the adequate answer that factually satisfies every condition of the question. In the initial stages, however, it is the absence of information that is the proper desideratum, the motivating cause, in that one continues investigating until one this vacuum is no longer apprehensible. It is not motivated by any achievable positive external datum, but by the desire to eliminate the consciousness of an information deficit, and this is true of an inquirer whether they indeed arrive at the true final answer and exhaust all relevant information or if they merely content themselves halfway there or even accept and believe falsehoods; it is true whether they are right or wrong that their ignorance has been satiated. "That's all I need (or want) to know" is not the same as "I know everything there is to know" but both are justifiable endpoints

of an investigation given that ignorance is the prevailing condition in any case.

The criteria of satisfaction are self-provided and determined by judgement. This obviates Meno's paradox, as the supposed final desideratum of inquiry as positive knowledge does not have to be obtained for certain. All that must be confirmed is that the irritation that isn one's self-consciousness of ignorance has subsided. This is true to the extent that belief formation and the satisfaction of doubt, as Peirce claimed, is the main goal of cognition rather than truth seeking. A truth seeker therefore just has higher standards for the revocation of their doubt, they postpone judgement and press on with their investigation well beyond the point that their own preconceptions have been gratified. Indeed Hintikka goes on to acknowledge that an inquiry is not necessarily a truth-seeking process. Information can be false, and the evaluative judgement of results is an independent working part separate from the process of information acquisition which itself only exhibits an eliciting function.

After Collingwood (1940), Hintikka finds credence in the notion that all questions stem from presuppositions that limit the scope of what may be asked, although he does take issue with aspects of the idea. A question always presumes a connection between its terms and imposes a structure which limits the range of available answers to it. For example, "What caused the dinosaurs to go extinct" implies that the dinosaurs went extinct, that dinosaurs existed in the first place, and perhaps if you are really unsparing, the concept of causality. No answer to this question which affirms the continued existence of dinosaurs, or denies them having existed , or denies they went extinct is permissible. That every question must call for a presupposition however makes little sense to me. (For what it's worth Hintikka seems to think that yes or no questions are presuppositionless because they are tautological and therefore logically empty.) To my thinking, for a presupposition to be the basis of a question it must be disbelieved. Unless there is an ulterior motive involved, for example one wishes to test what someone else knows by asking them a question and cares not for the answer, rarely are questions asked about which one already believes to be known. To hold a presupposition is to already assume one has the answer and takes it for granted. Presuppositions may be the result of answers from previous questions, in which

case one may have formed enough of a picture of the subject to guide further inquiry. At the initial stages however, it is precisely the absence of presuppositions which get things moving. I have already made it abundantly clear about what I believe triggers an inquiry.

Hintikka recognizes the problem of initiating an inquiry, and refers to Collingwood's assertion that all chains of presupposition trace back to an ultimate presupposition or initial premise. To me the content of this presupposition is obvious. It's merely the recognition "I don't know *this*." The use of the indexical "this" is essential because it picks out something specific. This is the most general form of a presupposition, from which any variety of more granular forms can be built out and refined as an information source is probed. Of course while "I don't know this" taken in isolation is nothing to work with, the recognition of it can initiate a process of speculation as to what actions might remedy it.

All told, I am willing to compromise with Hintikka and conclude that each question has a dual natured pairing of presuppositions. The assumptive nature of a question, which assumes the relations it expresses and imposes this form on the range of answers, has positive content. Yet there is also a negative presupposition at play which comes first, which serves as the setup function for the entire inquiry and which imposes its form on the positive presumption that flows from it.

In a similar vein is Hintikka's intriguing skepticism about Bayesian inference. Hintikka is suspicious that belief drives inquiry, placing it instead, after Peirce, at the end. The reason for this is that Bayesian reasoning presupposes an a priori selection problem of priors. To properly contextualize the input, Bayesian methods beg the question of how the right prior probabilities were chosen. There is, "…no presuppositionless Bayesian inference. Such an inference relies on a prior probability distinction that inevitably embodies assumptions concerning the world." (p. 201). Furthermore, to the extent that presumptive priors anchor the directionality of inference, circumscribing what counts as evidence for it or against it, there is always the chance that a chain of reasoning might be consistent with itself but wrong overall. "In the long evidential run, one's favorite measure

of information might very well turn out to be dead wrong." (ibid).

The glaring flaw here is Hintikka's neglect of the principle of maximum entropy and the isomorphic concept of the uninformed prior (AKA diffuse prior) used in Bayesian models. The uninformed prior is not exactly completely uniformed, it instead assumes the minimum of information as the prior probability distribution. It is often associated with Laplace's principle of indifference, which states in the absence of any evidence weighing way way or another, all possible conclusions should be assigned an equal probability. This forms the most "objective" starting point for a prior distribution without evidence. Initializing a Bayesian inference engine with an uninformed prior is akin to "looking at it with an open mind", of assuming the least about the data set and thereby avoid a biased model which would overly constrain the model. You might assume, after Quine, that this entails modeling the entire universe as that is the system with the maximum entropy, but in practice, the universe is never providable as a viable training input, and if it were, we would likely encounter the same problems as Wolpert's refutation of the feasibility of Laplacian Demon-like universal inference engine, the model itself being part of the universe. In any case the universe is unobservable in its totality, so it's an idle consideration to speculate further. The least is assumed about a data set for which assumptions can eventually be inferred.

One can find similarities in the use of a diffuse, uninformed prior with the notion of the intentional objectification of ignorance. Any inquiry by necessity begins with the minimum amount of information. What is the one piece of information of which the inquirer can be absolutely certain? That they don't know *something*, and that there must therefore be a something there is to know. From this vague, maximally ignorant initial state, more definite approximations may be sharpened down. The fact that the agent recognizes they don't know an inner something, immediately picks out an outer something to search for. It resolves the selection problem and orients the model toward a dataset that "fits." From this zeroed, null state, a more shapely characterization of the problem space can be developed. Another way of putting it is that while there is indeed no perfect rule to select a prior for the distribution, in the absence

of a good practical or empirical reason to pick one over another, the prior which assumes the least and which has the potential to capture the most possible information from observation, leaving nothing out because everything is initially admitted, is a solid fallback. It's akin to the idiom, "cast a wide net." This holds by the principle that if a solution resides somewhere between a minimum and maximum value, the assumption of an initial value closer to the maximum is more likely to encompass the solution somewhere under it than if the priors were under-specified and approximate the minimum.

Hintikka's critiques of Bayesian inference naturally flow into his critiques of cognitive bias, which he views as a wrongheaded idea. Indeed, he doubts their existence outright, going so far as to call it a "fad." In *Socratic Epistemology* he dedicates the entirety of chapter 9 to a takedown of the concept. He takes particular aim at Tvarsky and Kahneman's influential and original work on cognitive bias. Specifically he targets the conjunction fallacy as a crucial conceptual hinge on which cognitive bias depends and with the implication that debunking it also casts doubt on the remainder of the supposed biases.

In simplest terms, the conjunction fallacy is when two conditions are assumed to be more likely taken together than either taken separately. This is fallacious because the rules of probabilistic inference hold that in the absence of evidence, the assumption of a single condition is more likely than the conjunction of two or more, because this combines the uncertainties of each condition which self-evidently yields a less likely result than the uncertainty of a single condition.

The logical form of the conjunction fallacy is given as follows. Let's assume that S is a general description of the background evidence about B. Without explicitly stating either, S weakly implies C and D about B. Therefore which of these is more likely to be true of B:

1. C is true

2. C ∧ D is true

The fallacy comes from the notion that because S weakly implies both C and D, their conjunction must therefore be more probable even the though combined probabilities of two events cannot be greater than their independent probabilities. The more probabilities you add together, the more watered down

the likelihood of their conjunction, but the human mind often seems to think that because the description S of B each loosely fits each of these conditions, their conjunction must constitute the most likely assumption.

The most famous example of the conjunction fallacy is the so called Linda problem. Consider the following scenario:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Based on this information, judge the likelihood of the following statements:

A: Linda is a bank teller (T)

B: Linda is a bank teller and active in the feminist movement (T ∧ F)

Tversky and Kahnemen found that the vast majority of survey respondents judged option B ((T ∧ F)) to be more likely despite its seeming incongruence with the laws of probability. Furthermore, further experiments in which expert statisticians were given the same problem yielded similar results, which led the researchers to theorize that the mind had innate biases that skew the judgement of probability.

Hintikka's argument against the conjunction fallacy essentially rests on the claim that the apparent bias is attributable to the description S constituting only a single information source, and that if multiple information sources are provided, conjunction is valid. Essentially, it is the same reasoning behind interviewing multiple witnesses in a court case. Each testimony is its own probability distribution, and as it turns out taking the geometric mean of multiple combined probability distributions leads to a more accurate result (Clemen and Winkler, 1999)[https://faculty.fuqua.duke.edu/\~clemen/bio/Published%20Papers/28.CombiningDistributions-Clemen&Winkler-RA-99.pdf]. Hintikka summarizes, "In a nutshell, I might sum up my diagnosis of the conjunctive fallacy by saying that given suitable collateral evidence, certain messages are self-confirming in that they tend to increase our confidence in their source (p 217). Supposing that answers to A and B, T and (T ∧ F) come from two different sources. As Hintikka

explains, "The background information \[S] does not tell us whether T is likely to be true or not. If anything, the story about Linda makes it unlikely that she should be a bank teller… In contrast the information that F (…that Linda is active in the feminist movement) is made likely by the background information… for this reason (T & F) coming from a different source of information might have to be assigned rationally a higher degree of credibility than the answer T, coming from a different witness." (p 216). That is to say, on basis of the background information, the most likely proposition taken independently is F, that Linda is active in the feminist movement, on the evidence that she was deeply concerned about issues of discrimination and social justice, and was politically active. This proposition however is not offered as a third choice. Option (T ∧ F), taken in conjunction with (T) from another source has more credibility than (T) taken by itself. This is a result of the fact that the evaluation of a probability is enhanced by multiple information sources.

This conclusion is a consequence of three requirements, Hintikka asserts:

1. The need for someone to modify their priors in light of new evidence
2. The fact that such a modification often amounts to a reevaluation of the degrees of reliability ascribed to different sources of information
3. The fact that the evidence conveyed by an answer may affect the very priors in relation to which the probability of the answer is to be judged

Given that the conjunction fallacy seem to disappear when multiple sources of information are included, it would appear that the fallacy is a byproduct of how the information is presented and the response options are provided within the experimental design, bias within the test rather than the subject. Were the conjunction fallacy instead an innate fault of reasoning, then the inclusion of multiple information sources oughtn't make out the choice of option (T ∧ F) to be rational. Were the conjunction fallacy real, we would expect it to apply in both cases. It should stick for both the minimum and maximum information cases, Of course, on its own Hintikka's explanation should be tested empirically to see how much it water it truly holds, but he is right in the sense that the underlying

(un)reasoning involved in cognitive bias can be stripped out and analyzed purely in terms of the iron laws of logic from which people are claimed to deviate.

Implicitly, what Hintikka seems to be getting at here is that cognitive fallacies disappear when we consider how the agent interacts with information sources in a plausible, ecologically valid way. Yet is it enough to go on? If some judgements can be inherently fallacious, Hintikka asks, "How come? How can it be that humans, including trained specialists, commit such elementary howlers? What is the psychological mechanism that creates them? In spite of the bulk of the relevant literature, I have not found a theoretically satisfying explanation of it." (p. 219) Indeed how come?

I have endeavored to find a preliminary answer to this question myself. In short form, there is no absolute necessity that human cognition as it its generally exhibited must exhibit perfect rationality. In fact from a biological perspective, such behavior ought to be more the exception than the rule, given that most of the bulk of the brain's architecture is a holdover from an ancestral, animalistic pre-rational time before the capacity for abstract thought emerged. And people, all of us, can and do display stunning irrationality. The question is if there are innate causes of these distortions. People can behave irrationally even if provided all the information and variety of information sources in the world. Humans were not evolved to be precise and meticulous reasoners. Least of all were they built to be intuitive probabilistic thinkers. They were designed to survive and reproduce, to win arguments in order to gain social status, and even to deliberately misrepresent circumstances to themselves if the resulting misbelief happens to benefit them. The psychological mechanism is evolutionary. It stems from the brain being a kluge, a loosely assembled collection of improvised subsystems rather than a finely tuned and honed instrument of ratiocination. If evolution happens to favor the irrational it will be prioritized in the thinking process over and above the rational. Hintikka seems to think that people in general hold themselves to the same standards as a logician such as himself. This is no the case. There are good reasons that structural ignorance is implicit to the operations of the brain and mind. There are tacit sufficient reasons behind the unreasonable, a purpose underlying how people can behave with

seeming illogicality.

Ironically it deserves to be said that Hintikka himself falls for the fallacy of hasty generalization. On what grounds does he have to extend his -flawed- critique of the conjunction fallacy to all cognitive biases in general? He is subject to his own lack of conditionalization of his priors and an absence of multiple information sources he himself is critiquing?

Ultimately, Hintikka reduces cognitive fallacies to the Bayesian prior probability selection problem, which as I have pointed out by taking the prior which assumes the least about the data, namely the uninformative prior. Given also the faulty assumption that multiple information sources, cross-checked against each other, disproves the existence of bias, and given also that "how come?" can be in fact be answered, the concept of cognitive bias arguably still stands.

Adding everything up, Hintikka's analysis of inquiry, while insightful, sophisticated and rigorous, is still confined to the epistemic positivism I'm criticizing as insufficient. In spite of his correctness that information acquisition and question-asking is an important component of a more complete epistemology, he only goes half way and his approach is thwarted by ignorance blindness. Absent in his analysis is an understanding of ignorance which is indispensable to understand the context in which inquiry unfolds.

This concludes what I have to say about the theory of inquiry proper. While I make no claims of definitiveness, I am confident that the elements I have identified and the arguments I have articulated are essential to a fuller appreciation of the inner mechanics of inquiry. Before moving on to the last part of the book where I will address criticisms, tie off loose ends, and reflect on the human significance of these ideas, I would like to close Part III with an illustrative example of the theory of inquiry that's been sketched out. While inquisitive behavior is a universal human trait, there are certain areas of human endeavor which exemplify the logic of inquiry better than others thanks to their formalities and organized procedural nature. The scientific method in action is one such domain of systematized inquiry, as are congressional hearings, job interviews, and the like. But the one I would like to turn attention toward is the

judicial system. Namely, the role of the prosecutor in criminal trials. A criminal prosecutor is similar to a philosopher in some important ways. They aggressive pursue the truth, engaging in strategic questioning in order to isolate the unknown and corner it. Their job is also reconstructive, having to assemble from fragmented pieces of evidence a believable, and ideally one would hope a true argument. They are not just inquirers but interrogators, even inquisitors. What can the legal system tell us about the logic of investigation?

Mechanisms of Investigation

When someone does not know what they want to know, they have two choices. Resign themselves to ignorance, possibly fabricate some pseudo-knowledge that satisfies their desire, or turn their attention toward the outer world in search of some finding that brings their agitation to an end.

Another option is for someone who knows better to take note of the perplexed's cluelessness and help fill them in. If that was the only option we would be like the medieval scholastics who thought that Hippocrates and Aristotle were the apotheosis of knowledgeability. Such an attitude beckons a hopeless regression. They all needed to be told. One cannot pose questions in a language to something that doesn't speak it, and this fact alone shows the limit of what linguistic inquiry is capable of. Eventually all chains of epistemic dependency terminate at points where there was a discovery rather than a transmission of knowledge.

Who did Hippocrates learn from? Assuredly someone more knowledgeable than him. If that were the case, why not follow Hippocrates' teacher? Hippocrates happened to have come from a line of physicians, his grandfather being the first. Seemingly paradoxically, however, each one of the previous knew less than the next. If each previous learner simply trusted in their masters, the net result would be a stagnant recapitulation of previous knowledge rather than a steady expansion. Even if this pattern were to hold, there would have to be a progenitor who set the whole process going; human lineages, like the flow of

knowledge, do not extend back to infinity. The only way out then is inquiry, to break the cycle of teaching and to explore independently. To include what wasn't into what was known, and often to refute, augment, or revise what was in the process. If knowledge has grown it means someone went beyond the current state of knowledge and ventured into the outland reaches of ignorance.

In its most abstract form this problem of expansion of knowledge is related to how such a knowledge system is to see past itself without destroying itself. Ignorantly, present flawed knowledge seeks to at times ward off future improved knowledge, because nobody likes the feeling of a burst bubble. Nobody wants to find out what they thought they knew was the full story was merely the preface, a small piece of a much bigger truth, unless they happen to be intellectual disciplined and skilled at ignorance management. For finite knowledge to expand it must reject, shed, itself. Old knowledge must be seen as a mere placeholder for pending future knowledge. There are many reasons why knowledge did not drastically accelerate during the dark ages, (though admittedly this is a Eurocentric statement and at the same time the Islamic world was flourishing intellectually.) Yet one of the strongest reasons for the intellectual stagnation of dark age Europe is that the ancient Greeks, coming from what must have seemed like a fallen age of great civilization, looked to have everything understood in comparison to dark age scribe's present ignorance. Otherwise, how could their civilization shine so lustrously? So the scholastics set their sights on the target they could see, not on what those who they looked up to were targeting beyond their sight. Their chain of epistemic dependence became a manacle locking up their minds and caging their curiosity.

Only by becoming conscious of one's ignorance can more knowledge be actively created. Thankfully, in the best conditions humanity is well equipped for the task. Human beings have an emotion that causes them to move against their ignorance without strictly needing to identify it: *curiosity*. Such epistemic emotions, as they have been called in the literature, supply a raw boost of inquisitiveness that gets us over the initial meta-ignorance hump and kickstarts cognitive bootstrapping. All there needs to be is a generic and unspecified desire to understand the world or a given topic, just a small hint that there is something

more and a desire to find out where it leads. Human beings—at their best— seem to be genuinely drawn to the unknown and driven by a primordial desire to understand the world such that they unconsciously increase the odds of chance exposure to new discoveries by being curious. Those who learn are not those who know they know, but those who doubt they know. Implicit to the emotion of curiosity is a notional awareness of information incompleteness even if precisely what one is missing is not pictured by it.

Early on in life such exposure to the unknown is inescapable. Children typically have an exuberance of curiosity about their surroundings, indeed, it would be evolutionarily disastrous for them not to be. This curiosity can sometimes get them in danger, "curiosity killed the cat", without curiosity the cat —the saber toothed cat— would have killed them. Overall, curiosity must be beneficial or else we would have long gone extinct. If humanity did not find out about its world it would be consumed by the brutal natural of the elements. One of the few things we have going for us in an animal world otherwise ruled by the brute force of raw strength, speed, fangs and claws; by tigers and lions and bears who can physically destroy an unarmed human in moments, is our brains. The human brain is better conceived as an engine of doubt than knowledge. An incurious child is a sitting duck and may grow into an adult who is always running away rather than towards the unknown, fleeing from rather than conquering it, prisoner to assumptions which obviously did not solve the problem from which they flee. Children's abundant curiosity is a sign that evolution has produced a being that must be driven to learn without even even knowing what it should learn. Genetic programs beyond our comprehension give us a little push to start the learning process.

Stages of inquiry can be broken down into different phases ranging from the more general to the more specific. The raw curiosity of childhood ultimately leads to a more refined model of interests over the course of growing up, and the resulting segmentation of a subject into known and unknown parts bounded by an encompassing category. Generally the more definite one's awareness of what one doesn't know the more efficient learning. Education is a process for structuring an environment for the more ignorant among us so that they do not

have to spend their whole lives rediscovering the entirety of human knowledge. Eventually the student becomes more conscious of what it is they want to know, and they go off in pursuit of a career or whatever the case may be. Generation after generation, ignorance absorbs knowledge.

Natural curiosity eventually hardens into technique, solidified patterns of knowing, and perhaps inevitably science. It is a distortion of history to say that science was invented in early modern Europe in the Scientific Revolution. The general habits of science have been observed across a wide diversity of cultures including remote indigenous societies (Snively, G. and Corsiglia, J. (2001), Discovering indigenous science: Implications for science education. Sci. Ed., 85: 6-34. https://doi.org/10.1002/1098-237X(200101)85:1<6\::AID-SCE3>3.0.CO;2-R). At its most basic level science is an intrinsic human behavior, though not in an explicitly formalized methodology.

Science is not so much method of knowledge creation but an algorithm for coping with ignorance. To formulate and test a hypothesis is really to measure how wrong one is, not how right. It is motivated by a desire to tease out new data from experiment, not to simply confirm assumptions. Through a continual feedback of hypothesis and result and readjusted hypothesis the data is continuously better characterized. Directed, intentional testing presumes some description of the problem and cannot occur without establishing a context of errors around which to pursue it. Truth falls within the margins marked out by the errors. “I have not failed. I've just found 10,000 ways that won't work,” said Thomas Edison, and by failing so much, he bracketed the solution space within narrower bounds.

When we consider the explorers, inventors and discoverers of history we find that they often at first proceeded with a minimal overt understanding or anticipation of what they were going to discover. They were merely pointed in a direction, either in real space or thought space, that made the discovery increasingly more probable as they pursued it. What motivated them was a mere hint of something more or something else entirely that led them to make the discovery as accidental byproduct of another goal. Microbiologist Alexander Fleming discovered penicillin by accident upon returning after vacation and

finding a green mold growing in his lab that happened to destroy bacteria. To an extent any discovery must happen by accident, or come as a surprise, otherwise what was discovered would simply be a reminder of what one already knew and therefore not a discovery at all. Europe existed for centuries in ignorance of the two American continents to its west. When the age of exploration came to pass the plan was to find a more convenient sailing route to China and India. They were wrong because two whole continents were in the way. (Although of course this discovery was only relative to Europe.) But the wrong idea was enough to get boats sailing. For the same reason that a child cannot help but learn, people in the past had an easier time discovering new things. The set of undiscovered things is less exhausted at an earlier stage than it is at a later stage.

Often it takes many failed trials to lock in a discovery as the tester measures, misses the mark, remeasures, and misses again until the problem space is narrowed. Preliminary forays into the unknown are marred by numerous mistrials, errors, and miscalculations, but the initial work of developing the context for discovery needs those mistakes to define the liminal margins within which to progress. Errors define the minima and maxima of the correct result somewhere in the middle. Indigenous cultures have an in depth folk knowledge of the botanical properties of their plant ecosystems, knowing what is poisonous or what can ease a headache, only because this knowhow came at a steep cost in the consumption of many toxic plant substances over generations and the transmission of this information over the generations. Their knowledge was painstakingly paid for in a debt of stomachaches. This sort of eliminative trial and error, as we have seen, is the first recourse for overcoming meta-ignorance, those unknowns unknowns for which there is no prior insight with which to gain purchase of it and which can only come into awareness fortuitously.

In any case exploration always begins with an awareness of the unknown. The unknown is the cause of exploration. What exactly is unknown may be completely vague, but there is always a general sentiment that something novel lies “out there” egging us on. The explorer has an epiphany that what they once believed was everything is anything but and becomes aware of an ambiguous but enticing expanse beyond it. Such vague ambiguity is met with *curiosity*, an

instinctive interest in investigating what lies beyond one's ignorance without necessarily any other motive beyond what the emotion induces by itself.

Once curiosity sets in, next comes *probing* behavior, the hesitant, tentative poking around of the topic of interest in order to feel out its boundaries and coordinates, often through safe or restricted means. The waters are tested to see if it's advisable to continue deeper into the unknown. The objective of probing behavior is to sound out the context of the unknown to see where it fits in relation to present knowledge before committing to exploring it in greater depth.

The question arises: how does one come to understand pure novelty? If none of one's prior conceptions are a good model for some new stimulus it's a bit of a paradox how the purely unknown becomes understandable. *Ex nihilo nihil fit*, nothing can come of nothing. Of course, there is a fundamental dissimilarity between known and unknown, in that the unknown is not merely undetected knowledge "out there" waiting to be picked up, it is the potential for knowledge. The known is not like so many apples in a basket and the unknown is not apples on the tree that have yet to be picked. To understand the utterly novel unknown, we must think it through rather than engage in a sort of automatic information uptake. Novel discoveries require an investment cost of analysis and investigation to become knowable. At its most basic level probing is the act of interposing the known into the unknown so that a beachhead might be established into the latter and sense can be made of it in terms of prevailing preconceptions. Conjecture and hypothesis are ways to stab at the unknown using a finite knowledge base. Analogy is a well-studied example of a cognitive process that attempts to retrofit novel concepts to be in congruence with familiar ones (Similarity and Analogical Reasoning ,1989). Perhaps as the bible says, "What has been will be again, what has been done will be done again; there is nothing new under the sun." That a new discovery is made suggests that it already fits within the coordinates of the already known. By mechanisms of analogy, the unfamiliar is assimilated into the familiar. Unknowns which are so alien, so far afield from all that is known, are undetectable. The conditions of possibility for making them known have not been met. In the same way, a caveman cannot discover smithing because he has not yet discovered fire. In

Heidegger's terminology, such distant unknowns are undisclosed, outside of reason. Only when prior knowledge is built up to the appropriate level can it map to the truly novel unknown.

When contemporaries Newton and Leibniz independently discovered/invented calculus it was no mere coincidence. They were both brilliant mathematicians who presumably read the same texts and had a comparable understanding of all of mathematical knowledge up until that point in time. Math as a discipline had built up to such a point where calculus was the next latent evolution in its knowledge, and they were the most probable conduits to instigate this next stage. As Hume stated, the inductive method applies to situations where "instances of which we had no experience resemble those of which we have had experience" (THN, 89), then the impression of the unexperienced event must be prefigured in what is already familiar.

This line of reasoning that the resemblance of the unknown to the known is what makes it discoverable seemingly entails a regress problem. At some point there would be insufficient past experience to serve as a basis for mapping to the unknown. The key difference is between targeted discovery and open discovery. Prior knowledge allows one to hone in on a specific body of related ignorance congruent with it, but having the whole world before you be unknown also creates maximal conditions for incidental discovery. Early on, in both an individual's life or the history of the species, opportunities for targeted discoveries were minimal but for incidental discovery maximal. The fact of accidental discovery is sufficient to establish a "seed" of basic knowledge to refine into more directed explorations. When negative ignorance is highest learning is at its most probable, all things being equal, even if what one learns was by accident. (In an interesting contrast, when positive ignorance is highest learning is at its lowest or in a way goes negative, because wrong ideas can only lead to wrong conclusions, deepening ignorance.)

To get a better sense of the logic of inquiry, it's helpful to remember that human beings are not what Andy Clark calls "cognitive couch potatoes" passively awaiting inputs. Rather, cognitive behavior is *enactive*, always seeking out inputs to meet predictive signals in the middle. So it is with probing

behaviors. Whether it its sending a literal space probe into the solar system to take in data about the cosmos, a radar picking up the location of an object by bouncing radio waves off it, a bat using echolocation to hunt down prey, the act of sniffing a week old bottle of milk to detect if it's spoiled, or even tossing a rock down a dark well to sound out how far down the water is; in each case and many others, the unknown is actively engaged with in an attempt to minimize ignorance and manage uncertainty. Through enactive probing ignorance is aggressively encroached on and cornered. Effective epistemic agents do not wait for their ignorance to get the better of them, they actively scan for holes in their knowledge and accomodate that into their cognitive cycles.

Such investigational probing behavior reaches beyond body-bound sensorimotor activity and into "external representations" and external memory sources. These are experimental tools and detection equipment such as microscopes or telescopes, or information encoded props such as written reminders or notes, that allow people to work around the structural ignorance imposed on them by bounded natural senses or memory constraints by using technology that embodies knowledge or memory.

Consider that even a rock can be used to this end. To toss a rock down a dark well tells me a few things. The delay between dropping it and the splash carries information about the distance between the top of the well and the water level below, from which I can also derive the length of the well's shaft. The sound of the splash or lack thereof tells me whether there is water in the well. Because our senses alone are inadequate—I can't see down the well— the ignorance we are seeking to resolve is not only about the object of inquiry but also the shortcoming of the senses. In other words, the limitations of the senses impose a kind of structural ignorance in that I cannot discover anything with them that exceeds their bandwidths. The fact that the naked senses are inadequate to detect the object of inquiry causes probing behavior using information revealing artifacts. More to the point, an instrument such as a telescope is specifically designed to embody knowledge about how to extend the farthest bounds of the unaided eyesight. To bypass, that is, a limit of structural ignorance imposed on visual perception by the normal constrained workings of the eye. The idea that

cognition can be extended beyond the brain and into the body and environment is an active research program (see Andy Clark, David J Chalmers (January 1998). "The extended mind". Analysis. for the landmark paper that popularized the idea). In basic terms, extended cognition is the theory that we incorporate our bodies, external symbols, environments, and other people into our thinking and that the outcomes of this extended cognitive system are richer than those of the "naked brain" operating in isolation. To my knowledge, what is not discussed in detail within this body of research is an ignorance-first epistemological inversion of the extended mind thesis. Tools such as microscopes or computers do not so much as extend the mind as they do mitigate contingencies of ignorance. They do not take what is there and enhance it but remove a limitation. While these two operations, enhancing versus removing an impairment, might seem logical equivalents of each other, (in the sense that to gain something is to negate not having it), there is a telling distinction. A telescope will not be of as much use to someone with blurry vision. Access to the knowledge is not the same as possessing that knowledge. In a similar vein, Socrates famously shunned the written word, believing that it encouraged forgetfulness and claiming it provided the mere faded resemblance of second-hand knowledge which is only properly and authentically derived from direct experience and independent judgement. In the dialogue *Phaedrus*, Socrates is recorded to have said:

> [It] creates forgetfulness in the learners' souls, because they will not use their memories; they will trust to the external written characters and not remember of themselves. The technique which you have discovered is an aid not to memory, but to reminiscence, and you give your disciples not truth, but only the semblance of truth; they will be hearers of many things and will have learned nothing; they will appear to be omniscient and will generally know nothing; they will be tiresome company, having the show of wisdom without the reality.

Socrates is essentially putting the extended mind thesis to rest long before it was created. He is saying that we offload information not to enhance memory, but to forget it and then we look over our libraries or the internet and fool

ourselves into think this knowledge is in our possession. Extended mind theorists argue that the external cognitive artifact is not strictly a part of the mind, but is "actively recruited" or "looped in" when it becomes relevant to the problem solving process, like a traveler referring to a map when they get lost. Of course the map has not become a part of the traveler's mind proper, but the use of it temporarily scaffolds it. *They remember what they will forget* and thus that's why they carry the map. In this way, extended mind theory is explainable as *ignorance management* rather than as mental extension. This is not so much an extension of the mind as it is a planned out ordering of ignorance (as potential knowledge) in the course of problem solving. That is, the map-user knows that they haven't memorized the directions, and they know that what they haven't memorized (don't know) is represented by the map. So referring to the map becomes an act of ignorance management, in which need-to-know, known unknowns are placed within close reach so that when memory fails it can be called up. It is unnecessary to say the mind is extended and we can avoid some weird implications of theory if we just make the more sensible interpretation that what appears to be an extended process is actually just ignorance that has been judiciously managed.

This view on extended cognition offers a novel take on it that evades some of its most prominent criticisms. For example, the claim that the map in the above example constitutes a literal extension of the mind, risks invoking "cognitive bloat", a fuzzy definition of the mind with no outer bound that ultimately includes imprecise and counterintuitive "wide states." In other words, if the true cognitive process extends beyond the cranium, it becomes difficult to demarcate where the mind ends and the world begins. (Adams, F., & Aizawa, K. (2009). Why the mind is still in the head. In P. Robbins & M. Aydede (Eds.), The Cambridge handbook on situated cognition (pp. 78–95). Cambridge, MA: Cambridge University Press.) On the view presented here, wide states are unnecessary and are a confusion attributable to a vestige of overzealous positivistic epistemology. The wide state is explainable in terms of ignorance management, the epistemic agent knows what it doesn't know, and knows that it can work around this ignorance by filling its environment with external

information. The agent is able to setup an external information source to accomodate gaps in knowledge that it can expect to encounter in the course of problem solving. What are considered examples of extended cognition, therefore, are really reducible to an investigational process involving the self-recognition of expected uncertainty.

As a side note, reinterpreting the extended mind thesis in terms of ignorance management is a compelling case for how the theory of ignorance can recast traditional problems in epistemology or cognitive psychology in a revealing new light.

By using epistemic probing tools to investigate the unknown, the inquirer entangles what they do know with what they don't know, and so can infer information about the unknown by how the known is displaced by its interaction with it. When a radiologist uses a positron emission tomography (PET) scan to create a map of metabolic processes within organic tissues, they are engaging in just this sort of information entanglement. Injected radiotracers have a short-half life, and so eventually undergo beta decay that emits positrons (the antiparticle of electrons) that collides with the electrons in the atoms of the tissues. When antiparticles collide, they are destroyed and energy is released. The ensuing annihilation of both positron and electron causes the emission of gamma rays , which then strike a detector that converts this scatter pattern into a visual representation of the tissue. On a fundamental level this highly advanced scientific procedure is absolutely no different from the simple rock down the well example mentioned earlier. In both cases, something we can be sure about is interposed between something unknown, and the displacement affected by their interaction encodes some information.

(p 176)

When considering the object of inquiry, the extent to which it is unknown is also the extent to which the maximum of possibilities can be attributed to it. It's unclear what might be true of it, so anything goes and "your guess is as good as mine." The principle of maximum entropy states just that: the highest level of uncertainty is the situation in which all possible infostates are equiprobable. Probing is an attempt to dent this probability landscape so that it falls out of this

uniform distribution and various outcomes are observed to be more or less likely. So the process of investigation is all about establishing a bridge into the unknown using a probe that unbalances the even probability distribution of maximum entropy, to downgrade the likelihood of some outcomes and upgrade the likelihood of others based on how observation and object interact as the investigation unfolds. Simultaneously and in step with this process of sorting out the empirical distribution, the object of inquiry becomes more concretely characterized as what is supposed about it falls by the wayside and what is actual about it comes to light. What we call speculation is conjectures about an object of inquiry that are unable to penetrate the maximum entropy distribution for lack of evidence or access into the nature of the events under investigation. Investigation *disambiguates* (rules out arbitrary possibilities), it sorts out the various likelihoods of what may be factual and what is merely suppositional, showing us where the facts begin and assumptions end.

Scientific inquiry is the paradigm of structured investigation. We see that rather than being categorically different from the sorting of equiprobable entropic infostate distributions, it is a specialized application of this method that has established various epistemic procedures to ensure that the role of subjective assumption is minimized in the disambiguation of empirical phenomena. Science is often described as the engine of knowledge, but it is more accurately a method of ignorance management. To ask a question is to assume ignorance. To hypothesize is to implicitly project one's own uncertainty as to the deciding of the hypothesis, to defer judgement and to evaluate how the results of experiment, previously unknown, settle the hypothesis.

In the process of investigation what is of fundamental importance is the boundary between known and unknown. This line of demarcation is a highly active, protean and volatile area, in which the changing shape of what is known and found out shapes what remains to be known and vice versa as two coupled and coordinated domains. Like an army planning an invasion, the forces of inquiry engage in much staging, scaffolding and preparation on the frontline of the unknown. There is much *strategy* in how the unknown is conceptualized in relation to what is known so that we can ask pertinent questions about it.

Ultimately when this planned "invasion" is initiated, the prongs of the attack are *questions*. Like various pincer movements and flanking tactics, questions corral and corner unknowns, creating isolated pockets that are "surrounded" like a player in the game of Go surrounding the opponent's pieces with their pieces and thus claiming their "territory." While this militaristic metaphor is admittedly a bit on the nose, it's worth noting that to "attack a problem" is sometimes used in colloquial speech, and that it is not unusual to view a cognitive difficulty as a kind of enemy to be vanquished through tactics and strategy. It's natural to think of ignorance as an enemy to be routed and aggressively defeated. Questions, in this respect, are the weapons in the War against Ignorance. Consider a legal prosecutor who approaches the defendant's case, mounting an offensive assault on their position, "grilling" them with various inquisitional techniques in order to catch them in a lie or expose a chink in the armor of their defense. As we will see later, the way in which courts operate are actually an intriguing philosophical specimen for the theory of inquiry worthy of its own separate analysis.

What is a question? In asking this question about questions we are exposed to the very content that it both consists of and pursues as its meaning. By asking questions about questioning we draw out the content that we seek to examine within the very tool with which we examine it. Much like an epistemic probe, a question is an arrow of knowledge shot into the dark. Or to use a more accurate metaphor, a question is a bridge from ignorance to an information source. It is a sign of knowledge of ignorance, of the self-recognition of ignorance that seeks to resolve its own incompleteness. To pose a question is to disclose an opening; it is to open the space in which the answer resides. Without first asking a question, the answer is imperceptible. To ask a question is to initiate a dialectic of investigation with an information source, the resolution of which results in an answer. (This poses another meta-question: must all questions have an answer?) A question picks out a region of ignorance in which to stage an "invasion." A question is akin to a *vector* in mathematics, a quantity with both a magnitude and direction. A question is a vector in discovery space, the topology of the known in relation to the unknown. Its magnitude is its semantic content which identifies a region of ignorance, priming it for active inspection. It picks out the

"something" to which it is directed. Its direction is that area of discovery space it seeks to penetrate which has been singled out as a known unknown of priority. To ask questions is to stage cognitive machinery at a certain point along the frontline between knowledge and ignorance in order to "dent" it and proceed further into other areas of ignorance that have been disclosed by the current state of knowledge. Such a region of ignorance must be probed as a prerequisite to accessing the deeper unknown beyond it. After all certain ignorances must be incorporated into the knowledge base for deeper ignorance conditional on it to be uncovered.

When a complex problem is under examination a question is usually never asked in isolation. Rather, the question stems from a theory about the object of inquiry, which itself evolves as the investigation proceeds, and is part of an overall strategic sweep of questions which seek to characterize the problem from all angles. As questions are asked the underlying theory from which they originate is sharpened and becomes less ambiguous and more definite. As we will see, there are dynamic dependency chains involved with questioning, in which answers at one stage open up further questions conditional on it that otherwise wouldn't have been asked. The question is but one of a system that attacks the problem in a united front. When questions are asked as part of a system, they suggest an organized mobilization against targeted ignorance, rather than isolated skirmishes. In the face of infinite (or more precisely uncountable) ignorance, the best knowledge consists of knowing what questions to ask. Again, the intentional objectification of ignorance is the essence of the logic of inquiry. All deliberate inquiry begins from a standpoint of identifying a knowledge gap and devising a means to fill it.

Implicitly to recognize the need to ask a question, explore, or probe is to first take the measure of one's ignorance. This self-assessment of ignorance is the first necessary step for motivated learning. Otherwise, if one does not see it as a cause of alleviation for one's ignorance, the external information source will be cognitively inert; it is not recognized at all as an information source. It is only when the information source is pinpointed as a possible solution to the problem of ignorance that prompted the question that it becomes a live possibility. The

question leads to the information source. Without the question to prompt it and bring it forward into awareness, it is merely an undistinguished aspect of the background scenery. To ask a question is to *activate* an information source, to recognize it as relevant and set it apart.

Whether mundane or extraordinary, whether trailblazing or well-trodden, exploration is the cornerstone of epistemic behavior. Both as a species as a whole and as an individual every time some new vista has been revealed by our actions, the vast and vague essence of the unknown, *res ignorans*, ignorance stuff, is gradually assimilated into *res cogitans*. The border between known and unknown is elastic and irregular, constantly redrawn such that some unknowns become more knowable and accessible than others and some less so. When someone colloquially says that such and such is ahead of their time, it can only mean that they have scouted deep into discovery space, creating an indent into which only their minds have ventured and which outsteps the most probable adjacent implications of present knowledge. They have "pushed the envelope." One way to think of genius is someone who is capable of "skipping steps" in discovery space in this manner. Arguably had he not skipped the steps, Einstein's views of gravity would have been reached by the natural drift of scientific consensus some indeterminate years in the future. In contrast when someone says so and so put us back several years, that society has slipped into a dark age, the border of knowledge is analogized as linear time, as a kind of slider with a pointer that can be moved back and forth depending on the circumstances. Such a regression in knowledge however happens not in time but in discovery space. The frontier of knowledge has been rolled back and rather than moving deeper into ignorance it goes backwards to the familiar and ignorance reclaims territory.

The variety of mechanisms that fascinate ignorance reduction are many. There is, however, one such method which stands above the rest and which has changed history utterly. I speak of course of none other than the scientific method. How does the logic of the scientific method manage ignorance?

Philosophy of science usually typifies two distinct approaches to the scientific process. The first is scientific induction, in which a lawful regularity is inferred

from repeated observational confirmations of experimental predictions. A scientific theory emerges from the predictable and statistically regular order observed by multiple experiments. A pattern is read from the data. Then there is the hypothetical-deductive method, which is sort of an inversion of the inductive method. Science begins with hypotheses, which can simply be creative guesses or speculative intuitions and then experiments are devised which test the hypothesis. Hypotheses that are not aligned with experimental results are rejected or revised and the cycle repeats. (A third type of inference, the abductive, has been proposed, but it presumes a theory from which one evaluates an observation as an "inference to the best explanation." We are more concerned here with how results are acquired, rather than how preexisting information is utilized to generate hypotheses.)

There is some debate as to which of these methods is the one scientists by and large embrace. Popper, in *The Logic of Scientific Discovery* sided with hypothetico-deductionism, although inductionism is generally thought of as more "pure" and objective. Popper believed that induction was inadequate to explain the essentially creative aspects of science, and argued that it was impossible to derive a theory from a corpus of evidence without presupposing an assumption. Both or one of these forms of investigation may have applicability, but the concern here is not primarily which approach science favors, but rather how scientific inquiry is structured to minimize uncertainty and eliminate excess possibilities in the course of deriving its results. "Results" here is used in the broadest sense to include not just data and research findings but the theories and explanations that relate and connect the data in a systematic knowledge system that informs all scientific assumptions.

At its most basic, the scientific method is simply a trial and error elimination technique, however it has much more on top of it. A scientific experiment tests the relationship between variables. Such tests are a measure of the degree of wrongness rather than rightness, because there is no notion that positive results prove a causal relationship. Positive results are then the least wrong results, though it is presumed that further experiment might reflect differently on it.

The scientific process can be likened to a buffer against uncertainty. It is

organized to carefully erect conceptual barriers and dams, to exclude as much noise and confounding variables from the experimental conditions as humanly possible. The precise deployment of dependent and independent variables is meant to shut out the possibility for uncontrolled variation and isolate exact relationships of cause and effect. The goal is to simplify to such a degree that only a single channel is open between the chosen variables. Granted even looking at it from the most basic physical level and disregarding all other sources of bias, it is impossible to perfectly isolate variables; even a force as meager as gravitational waves from distant stars or cosmic rays impinge on the most well-crafted experiments. The expectation is that a well designed experiment will isolate the variables to a great extent and interference reduced to a bare minimum.

What remains for even the most airtight experimental setups is still the lingering uncertainty of the best fitting explanation for the data. There are many infamous instances where a superficially fitting explanation, phlogiston, the ether, corpuscles, were levied to explain a phenomenon, only to later be overturned and replaced by a better explanation. Importantly, this process of explanatory formation and replacement can never be assumed to end in science —science is always pending, it is extremely epistemically cautious and reserved. Unlike a criminal investigation in the previous example, science does not inexorably tend toward a final state. Fundamentally it is a fluid, open and receptive process that constantly revises its hypothesis and lets nature do the talking.

Science does not trade in certainties and this causes one to wonder why such slogans as "believe in science" or "trust the science," are so prevalent. Here, what is worthy of belief is the method, though all scientific results are best viewed with doubt and the method itself is driven by skepticism. Scientists themselves do not "trust the science" because they understand that the data is not a definitive record and nothing is proven. Tacitly the scientific method recognizes uncertainty. It is in "perpetual beta" always iterating and only through some process of accretion does a collection of data and evidence congeal into powerful explanatory constructs that are scientific theories. Theories, however, are not

invincible and as any science watcher knows the textbooks of yesteryear bare little resemblance of those of today. If anything error is a more important signal that experimental confirmation, as this ought to transmit more information than the that which supports a previously established hypothesis. Negative results are just as important, if not more so, then positive results. Both restructure the discovery space and carve up the available possibilities. When we look at the history of science we see at the successive elimination of various candidate theories a the gradual narrowing of survivors in a process that vaguely resembles natural selection. Somewhat paradoxically, the more science progresses the less it perceives as each advance rules out a class of hypotheses that are inconsistent with the accumulating corpus of evidence.

The scientific method enumerates what isn't as much as what is. Null hypotheses counterbalance the hypothesis, suggesting what would be the case if it were untrue. Negative results are an often neglected but critical aspect of scientific practice because it helps sharpen and clarify the problem space by eliminating dead ends. It is to the scientific community's detriment that it tends to favor publishing positive over negative results ([https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3917235/]).

Positivity bias reflects the understandable sentiment that it is more prestigious and glorious to be right than wrong. That human drive for recognition gets in the way of a clearer picture of the actual data. One study found that up to 85% of experimental reports tend to record positive results, a suspiciously skewed figure ([https://link.springer.com/article/10.1007/s11192-011-0494-7]). Such bias toward positive results is indicative of the human weaknesses and sociological or institutional pressures behind the science, and is a revealing of case of organizational strategic ignorance. The pressure to acquire research grants also means impressing those with the capacity to dole them out with successful research and flashy discoveries. Scientists therefore have an incentive simply to not publish their negative results so as to feign the appearance of being a potent discoverer. But finding negative results *is* a discovery. To suppress it is to spite the fact that negative results are the reciprocal and essential complement to positive results and convey useful information. Without well documented

negative results, inefficiencies and redundancies may flourish and the result is wasted resources, time and effort, and a scientific process that is not nearly as rational as it is made out to be.

One way to think about it is a scientific theory is everything it isn't. It's the surviving remnant of a prolonged onslaught of criticisms and attempted refutations. Negative results subtract from the expected margin of error for a hypothesis by narrowing down the possible viable moves that can be made in the conduct of experiment. Such a trial and error approach is analogous to saying to yourself, "I have checked every room in my house for my keys except one." A countervailing positive force meets this negative tendency in the middle of course, but it's misleading to assume that science proceeds without the negative and only through brilliant acts of straightforward discovery.

Science on this interpretation is method of structured uncertainty reduction. It deals more in ignorance than knowledge. Ultimately science is an ignorance-first epistemic practice.

The Philosophy of Courts

Because practical knowledge, knowing how, is of great interest it's valuable to examine real world situations where the cloud of ignorance intrudes on our goals and where thought is always fringed with uncertainty. It's beneficial therefore to look at cases where the struggle against ignorance is especially fraught with high stakes. While it is true that nearly every undertaking of importance confronts uncertainty to some degree, some pursuits are unusually thick with it. Such pursuits often culminate in a single irreversible decision which decides the fate of someone or something. As an example of such a problem space involving high stakes inquiry the proceedings of a legal trial is a distinctive specimen, particularly the role of prosecution in determining the guilt of the accused. Prosecutors must use evidence, often fragmentary, and combine that with the letter of the law and cogent arguments to navigate through an adversarial information environment and locate the criminal in the act somewhere in the past.

Most nations use one of two types of trial convention, the adversarial or the

inquisitorial. The adversarial system is one in which a neutral judge and jury evaluates both the defense and the prosecution's cases and on the basis of the evidence and accounts presented decides the ruling. The inquisitional system instead involves the judging body actively interrogating the plaintiff and the accused and their representatives and taking on a more engaged arbitrative role in the conduct of the trial. These differences are however minor implementation details when considering that no matter the approach, the same purpose of investigating the substance of the case is at work, and the defense is either guilty or innocent. However the adversarial system is of greater interest philosophically, because it is a case study in which the information sources are at odds and conflictive. It is not a perfectly transparent investigational environment. Unlike science, in which everyone is incentivized to freely and willingly share data honestly and everybody benefits from open exchange, an adversarial court case is the reverse. The prosecution argues for the case of guilt and the defense the case for innocence. A zero-sum criminal trial is a struggle of clashing arguments and dissonant versions of events, a back and forth of epistemic sparring intended to induce ignorance of the truth or the knowledge of the truth in the jury, or to wedge the opponent into a position of revealing a lie or fumbling their argument. Winning is often easier if the truth is on your side, but even so, in some cases it is not so cut and dry and it certainly happens that a guilty defendant is successfully acquitted thanks to the finesse of a skilled attorney or the innocent is convicted regardless of whether it was a mistrial.

This gritty antagonistic argumentation to my mind reflects a realistic fact about many forms of inquiry which Hintikka's more idealized analysis overlooks. Which is that one cannot always simply reach out and take answers from a reliable, cooperative information source. Not all information sources are presented in good faith. There are resistant information sources and misinformation sources. Especially if that source is another person with something to hide, or who is deeply invested in preventing the truth from becoming known or in misleading the person asking the questions. Methods of inquiry must be at their most rigorous when you can't simply tap a few buttons, do a Google search, and get an answer or ask a well meaning teacher, but instead

must work around and outwit various epistemic hazards and pitfalls. Some answerers will seek not to make you wiser, but more ignorant, yet you must proceed to ask questions and bait them into revealing information. If not to get the truth out of them then to force them to answer in a way that is incongruous with it and thus exposes a weakness in their argument. This adversarial aspect to working around an unreliable information source in the process of inquiry is an important avenue for the study of inquiry.

In the process of any trial there are initial conditions that set the stage for the proceedings. Though they are far from perfectly rational, courts can be conceived of as rule-regulated state machines that compute toward an end result—the judgement— through a series of intermediate steps. The language of the law seeks above all to eliminate any ambiguity. The court system itself is regimented into a strict procedural device. This is a typical pattern when order must be maintained and the output of human behavior must be collectively directed and pointed at a single determinate outcome. A trial is prompted by an accusation, a defendant arguing for their innocence, and a plaintiff seeking remediation. Organized back and forth deliberations that obey strict regulations ensue. Supplied to the court is the evidence materials procured by police investigation, witness reports, and other relevant details that comprise a public information background to work against. The task of the prosecutor is to assemble these materials, connect them together in an argument and demonstrate beyond a reasonable doubt that the defendant is guilty as charged.

In many ways the court process resembles philosophical work and the similarity is no coincidence. In both cases and indeed in most cases where heavy argument and critical thinking must be endured or deployed, there is an effort to diminish unclarity, to collapse the possibilities of the unknown into a truth formula and wrangle uncertainty into a manageable shape. Both are heavily argument focused. With prosecutorial legal proceedings there are many variables and shades of difference that make it difficult to generalize about in absolute terms. That said in most cases the guilt of the accused is held in some doubt, innocence is presumed, and the onus is on the prosecution to disprove the case for their innocence rather than vice versa for the defense team. The events of the

crime, like all things in the past, has ebbed away, leaving only fragmentary hints and clues as to the nature of what really happened. To complicate things further the defense, assuming their client is guilty, is incentivized to dissimulate and evade, filling the proceedings with complex ignorances, and rendering the epistemological environment of the court all the more fuzzier by counteracting the clarity the prosecution seeks to achieve.

The important features of the prosecutor's task are essentially *reconstructive* — they must rebuild an accurate semblance of the events of the crime, or at least a convincing enough portrayal to suffice as proof of guilt. By "reconstructive" it is meant that the prosecutor must construct an image of the past from often deteriorated and spotty evidence, they must assemble a picture of a state of affairs whose singularness and coherence has since subsided into obscurity and lost in time. In the moment the event transpired, its fact was singular and coherent, it was a configuration of states with definite relations which then bifurcated into casually related but temporally dispersed effectual displacements. The events happened, and the world was shaken up a bit because of it, not quite the same as it was before. The states that composed the events fell out of coordination with each other like scattered billiards, and whatever evidence of their past coherence is in whatever "fingerprint" literal or otherwise left over; whatever indent or causal displacement of effects they impressed on the rest of the world in the course of transpiring. By retroactively connecting the dots, the prosecution reconstructs the now entropically diffuse remnants of the incident, the murder weapon, the DNA, the suspiciously open window. They seek to arrive at a depiction of the now subsided configuration of states and relations. A set of states that can now only be inferred from the faded impression they have left in a partial evidential record.

Any investigation about the past necessarily takes this reconstructive form. Any phenomena that aren't immediately occurrent must be reconstructed from the past or extrapolated from the future, but in this case the future has absolutely no role to play in the determination of criminal culpability. The proceedings is one of retroactive inference, or back-propagation, which goes backward from temporally antecedent effects to preceding causes, and through careful argument

excludes counterhypotheses that might falsely explain the antecedents. It is a paradigm of Bayesian reasoning. A backdoor left open at the scene of a crime may have merely been a coincidence, but taken together with other data—shoe prints in the backyard matching the defendant's boots, a cut in a fence, a toppled trash bin, a strand of matching hair— a constellation of related facts revises the probability of guilt upwards.

There is often little clarity at first. Initially (and ideally) the jury perceives the case as an undifferentiated and homogeneous spread of ambiguity. Their priors, ideally, are in a state of maximum entropy. The prosecution's task is to sculpt that ambiguity into a honed system of facts. Data point by data point, argument by argument, the jury's ambiguity is pared down until a decision is made.

The object of an investigation in a manner of speaking is not the event in question itself but the rest of the world that it has effectually displaced and left a mark. Evidence is each relevant displacement, that forms a ring our perimeter around the event like a cookie cutter outline. That things happened in such and such a way points to the world subsequently having a certain deformation that corresponds to the event. Such a "concavity" of effects etched into the world opens an explanatory hole for which only the one true cause fits.

The legal system is of course deeply human, tinged with emotionality and moral sentiments that both bend and transcend the strictures of rigid logic. Beneath it all is still this logical skeleton which informs the proceedings. Far from ideal scientific circumstances for the discernment of truth, criminal investigations nevertheless display the real-world, scrappy and imperfect conditions in which practical inquiry unfolds.

The End or Beginning of Science?

Is there an end to knowledge? In the book The End of Science, science journalist John Horgan argues that *all* or nearly all the fundamental theories of nature have been discovered. What remains then is to unify them under a theory of everything (more on this later), fill in the details and add some commentary. While there may be some nitpicks and controversies left over the ultimate frontiers of scientific knowledge have reached their maximal extents.

Accordingly scientific knowledge acquisition will be progressively more difficult and breakthroughs ever rarer.

In the spirit of thermodynamic metaphors, the situation of modern science can be likened to a thermal gradient that has reached near equilibrium. When the gradient is high on one side and low on the other, when, that is, science is in its early stages, it's easy to fit more "billiard balls" into the low density side of the gradient. As a matter of fact they flow effortless toward that side. As we learn more, the gradient evens out on both sides, and fitting more billiard balls becomes more energy intensive and impractical. Just as when early scientists had an entire vista before them to discover, modern scientists have narrower opening in which to make discoveries. They are almost at equilibrium.

The idea that science is exhausted is a deeply mistaken and overly bold claim given how it fails to answer the pitfalls of the metaknowledge problem (how we know we know and how to know there is no more to be known.) One piece of the argument is that existing knowledge in different areas of science perfectly and seamlessly align. If you exclude as of yet scientifically unconquered mysteries such as consciousness, the corners of all known scientific theories line up seamlessly, snapping together without any leftover nomological danglers or unaccounted for assumptions. Science is consilient. Certainly there are some contingent unknowns, undiscovered species of organism or particles or relationships or details between this and that theory, but the known laws and explanations map these new discoveries and none of these data points are inexplicable. Our ignorance of such novelties reflects an incompleteness of data not of theory. Any new and forthcoming data should neatly slot into an existing theory. That theory is in turn supported by the rest of the body of scientific theories, the background scaffolding of conceptual dependencies, psychology on biology, biology on chemistry, chemistry on physics…) that depart from any one area into other fields of scientific knowledge all of which hangs together as a cohesive hierarchical whole.

There is a flaw with this view that because everything lines up, you must have considered everything. Just because everything we know can appear to explain everything we're aware of doesn't entail we're aware of everything.

Indeed a false satisfaction may result from such thinking, a kind of naivety that lulls one into an epistemic dogmatic slumber.

Then there is this much hyped Theory of Everything the implications of which must be better understood. An actual ToE would constitute an inflection point, a singularity at which new vistas of thought should be inferred which cannot be accessed from the perspective of any subcomponent of the theory. In such a way then a ToE is only a beginning and as such is self-contradictory.

The truth is we don't know what a complete ToE would look like, or how it would lead us to as of yet incomprehensible insights. So even the achievement of a ToE would not satisfy the criteria of epistemological completeness. How could we know that it literally explains everything, that every fact is derivable from it, without first having known that we have experienced everything? This goes back to the problem of induction. For one thing the obtainment of a ToE could—and there is no good reason to exclude this option—open up the possibility of deeper regions of exploration not yet realizable from within the context of present knowledge. It may be a theory of everything with respect to everything that we presently know. But once obtained, it may reveal that we know a much less than once believed .

You can take it a step further. A putative theory of everything must also be, and I mean this in complete seriousness, a theory of nothing. It must resolve the set-theoretic contradiction that everything contains nothing, the negation of everything. These are precisely the laws of ignorance as described herein. A ToE will only ramify and clarify the fact that it isn't quite what it claims to be. That the more we approach a unification of physics and the other sciences, the better insight we obtain into what phenomena that unified theory fails to illuminate.

A strange thing presents itself here, known adjacently in the psychological literature as the Dunning Kruger effect. Low skilled individuals tend to misjudge their skill level, while higher skilled individuals are inclined to accurately appraise their limitations. A generalization of this effect could be ventured. Stated as such: the more you know the more you realize you don't know. Conversely, the less you know the less you realize you don't know. All of this adds up to say that a ToE is anything but. The achievement of it articulates the

terms of its own negation.

Of course what physicists mean by a ToE is not literally an explanation of everything. It won't tell you what will happen next Tuesday in your personal life or who will win the next election or what stocks to invest in. It does follow however, that if we maintain a commitment to the unity of the sciences, a synthesis of all physical facts opens the way forward to a greater harmonization of all empirical knowledge.

Those of us sympathetic to the "radical" brand of empiricism advocated for by James, resist this chimerical drive to completeness of the sciences. Radical empiricists hold that experience is inexhaustible. New relations between experiences jump out and re-implicate past experiences, establishing new grounds for inference and broadened investigational frontiers. The push for a ToE is the culmination of the positivistic, "inside-out" knowledge-first view of epistemology which we here counter with a negativistic, "outside-in" ignorance-first paradigm.

A ToE is more like a beginning than the ending of science. In a manner of speaking all of science up to it is a prelude. If one could derive biology from chemistry alone, who knows what new facts could be learned about biology which are simply unrecognizable to current science. To the extent to which a final theory is best approximated, the more aware of our ignorance we will become. According to my formulation the best and most useful knowledge is of this negative character. The struggle against ignorance in which our active experience is thrown is unending. So the most we can hope for is sharper clarity into what remains unknown. If nothing else a genuine ToE will provide more definite brackets around what it falls short of explaining.

Sometimes the way to progress is to retrace your steps back to where you started and to take a different path from the beginning. With respect to epistemology, it turns back to Socrates's famous dictum, "I know I know nothing." The statement almost carries a hint of cheekiness, in being both insightful and self-refuting. Socrates is both stating a truth about his philosophical standpoint and wittingly uttering a paradoxical absurdity. Clearly

if you know you know nothing, you at least know one thing. Knowing that you know x, is a separate fact from knowing x independently. Socrates therefore knows and does not know x. If correct, Socrates actually knows at least two things. He knows nothing and he knows he knows nothing. The second bit of knowledge however contradicts the first: he knows something in the sense that he has a definite grasp on the extent of his own knowledge, which is nothing. If that is the case then, how can he claim to know nothing? And on it goes. Such a maxim encapsulates the fundamental role of meta-knowledge in regulating the understanding of ignorance. Without knowing what we don't know, ignorance overtakes us and oblivion beckons.

Socrates's witticism epitomizes his radical skepticism. Ever restless of mind, questioning everything, doubting everything, never satisfied with what is apparent or taken for granted— Socrates recognized that the best knowledge is forever deferred, that to arrive at knowledge was more like to clasping one's hands around a puff of smoke rather than to clutch a solid rock. Deferring final judgement on the ascription of facts takes discipline but it is the best way to keep learning. For Socrates, knowledge was a process rather than a static consequence.

Socrates was the first to codify the practice of intentionally objectifying his own ignorance. He was aware, far before he had any right to be, that the system of human knowledge was in the germ in his time. His awareness of his ignorance drove him to ever more intense bouts of inquiry, causing him to become such a disruptive force in Athenian society that the authorities eventually had to pressure him into suicide lest his technique grew too popular and their power was openly challenged.

Socrates' discovery of the knowledge of ignorance came at the beginning of philosophy, and fittingly, it metaphorically resembles a seed. A singular kernel of an idea which was long dormant in the soil and has been a long time sprouting. Philosophy over the course of history embraced his critical mentality, but they forgot about what had been left buried, the roots of this critical attitude. Ignorance was an unworkable concept or it was trivial; the conceptual tools were not yet available to make sense of it without sounding contradictory, or it was simply unimportant. Yet the true source of that philosophical skepticism was

Socrates's radical acceptance of the absolute reality of his ignorance.

The wisdom of this acceptance of ignorance cannot be understated. Suppose that Socrates is in fact ignorant of the true state of human knowledge. He cannot be more disadvantaged by believing he knows nothing than if he believed he knew everything. The true state of his knowledge ,it is safe to assume, leans more towards approaching nothing rather than everything. Any finite quantity is closer to zero than it is infinity. The belief that one knows next to nothing leaves open the possibility for knowing more, while the belief that one knows all there is to know is the first sign of senescence, of cognitive decline that slips ever into a deeper and more terrifying ignorance--ignorance that is irreversible and uncorrectable. Even if one knows much, the mental habit to doggedly fixate on one's ignorance over one's knowledge is an invitation to greater knowledge. Even the most knowledgeable experts ought to be on guard against lulling themselves into false certainty. In zen buddhism this habit of fresh-minded openness and receptivity to learning at any stage of training is called *Shoshin*. "beginner's mind," (Shunryu Suzuki (2011). Zen Mind, Beginner's Mind. Shambhala Publications. ISBN 978-1-59030-849-3.) This beginner's growth mindset is what Mark Zuckerberg partly credits for his wild success in business ([http://change.is/quote/always-keep-a-beginners-mind]). For even if they know an exceeding amount about their area of expertise, it's important also to keep in mind the relational nature of knowledge. Everything is connected; knowing all about *this* does not tell you how it relates to *that*, and to understand the combination of these factors requires new synthetic knowledge not represented by domain-specific expertise.

If the amount that can be known about the universe in absolute terms is infinite or exceeds the memory capacity of the human brain, surely whatever one knows is closer to zero than to infinity or to an amount of information which cannot be physiologically supported. Therefore the most intelligent and rational conclusion is to fess up and admit to one's ignorance. It's the path to growth.

Someone might claim that the views I advance are pessimistic. *Ignoramus et Ignorabimu*s, "we do not know and will not know", said German scientist Emil du Bois-Reymond, expressing a belief of ultimate epistemic horizons that does

not one to one overlap with all possible facts, specifically those regarding the true nature of physical mass and force, the origin of motion and the nature of consciousness. Of course as a statement of fact the natural response to this is: how do you know that? How can you generalize from your own personal ignorance to the conclusion that what is not known now can never be known? Elsewhere we went into some depth to demonstrate that knowledge of what is in-principle unknowable is tenuous at best or impossible at worst, because to affirm knowledge of the unknown is to reference in terms of what is already understood. What we can claim to be unknowable are certain somewhat trivial meta-propositions about what is knowable. The unknowable cannot be recognized by knowledge because it shares no potential connections to any possible knowledge graph. As argued in chapter 3, the linking of a known to an unknown is a preliminary step toward integrating the unknown into the known. knowledge is a link that connects facts. The very fact of recognizing the unknown means that one can formulate an idea of what that unknown is, it is *comprehensible*. What's more, to draw the line and claim something is unknowable in absolute terms necessarily presumes knowledge of all the pertinent facts and the knowledge that one knows all the pertinent facts, a patently unsatisfiable condition. Therefore we can only say with confidence that it is in principle unknowable what is unknowable beyond this one indexing meta-statement.

Ignorance casts a long shadow over comparatively puny knowledge. The abundance of the unknown overwhelms the relative poverty of knowledge. This knowledge of the unknown is easily achieved: I can not tell you the position and velocity of all the immense numbers of particles in the universe, and I can't tell you what they'll be tomorrow or at any time forever: thus there is an astoundingly large baseline set of admittedly trivial unknowns. One might say, "should we not be proud of our hard won science, should we not be confident in the indisputable problem solving powers of the human intellect? Powers which have overcome staggering adversity over the course of history and solved the most astounding riddles? You yourself say that much of what is unknown comprises trivia, the location of a dust particle a million lightyears away, the

cause of death of some nameless and unnoticed insect, the source of indigestion in a random field mouse's gut. And you say that knowledge of general principles allows us to compensate for ignorance of the particulars by abstracting their properties into universal category prototypes or statistical models. So why should we surrender to ignorance?" And I would respond: absolutely. Human knowledge is indeed extraordinarily powerful and perhaps even unlimited in its potential, but it will never catch up to the totality of ignorance. But that doesn't matter. What matters is whether all *interesting facts are knowable.* knowledge must have some ignorance to work with to keep growing, and this is the first and last sign of its perpetual immaturity, its permanent condition of "perpetual beta." Problem solving has not gone out of style the more we have learned, and problem solving is fundamentally the act of mobilizing knowledge against a targeted chunk of goal-inhibiting ignorance. Have we stopped having problems? Does anybody truly believe we ever will? Do we not still encounter surprises daily, weekly, monthly and yearly? If the answer to many of these questions is yes then we can be sure that ignorance will always be one step ahead of us. This is not a depressing surrender to futility but a plain common sense acknowledgement of a neutral fact of nature. Indeed the more knowledge is acquired, the more layers of reality are revealed and old problems are replaced by deeper, newer ones.

On the contrary the view of ignorance advocated here is optimistic and empowering. It shows that ignorance is lawfully organized and not some inscrutable bewitchment. We can say with confidence that ignorance is a known known. It shows how its very logical boundaries can be leveraged to encircle and fence it ignorance rather than permit it to run wild. It's true that the deepest ignorance is that which cannot be bounded by any knowledge, the deep unknown unknown, and it is this kind which will without fail blindside us at the most inopportune moments for as long as we live. Ignorance is such that it has often done its damage by the time it is too late. Retrospective knowledge is often humiliating or insulting: how could the obvious escape us? We can however guard ourselves against the threat of the inevitable unknown unknown by enhancing our sensitivity to the prospect of our omnipresent ignorance. And

when life humbles us with those stinging hindsight reminders of what we didn't but should've known, there is no need to blame ourselves too much over it. Instead we can remind ourselves that such an outcome is actually the most probable given the vastly disproportionate ratio of known to unknown that leans well in favor of the latter. The universe, we can often say, made me not know. knowledge is the exception to the rule, but this also means it is exceptional in the other sense of the word: remarkable.

I am under no illusions that the goal of completing epistemology has been achieved. Somewhat gratifyingly, this confession of inadequacy is consistent with my premises. One can never really say, "and that is all" in good faith once one has compassed the true scope of ignorance. The situation ever evolves with ones relation to ignorance, both on a personal level and in its study; new findings emerge and old ones wither away. The whiff of potential new discovery always clings near to the air, and it is always possible to take another step forward deeper into discovery space.

Astonishingly, or perhaps not so surprisingly after all that's been said, there's always more waiting. Of all shortcomings of this book the most glaring is that it is the traditional branch of epistemology and its trove of basic arguments which I have left out from this account. It has been left mostly assumed, preferring to treat it as a kind of pre-defined term in the triadic combination of knowledge, ignorance, and inquiry which is the full picture. And it posits that it is the latter two terms which need to be brought up to the same high standard of philosophical respectability already enjoyed by the theory of knowledge. An entirely separate project could be made out of systematically reviewing the established state of the art in conventional epistemology to demonstrate how an ignorance-first perspective illuminates its darkened corners. This would indeed be a valuable exercise, especially if the theory of ignorance is to prove its usefulness and put weight behind my claims that its inclusion clarifies longstanding epistemological confusions, sorts out deadlocks, fruitfully reinterprets various puzzles, and resolves controversies. While I have touched on it here, the bulk of the work remains to be done. In a very real sense this is the actual task and what I have articulated is merely a preface for it. Nevertheless I

am confident that the appropriate framework for this pending task has been substantiated in these pages. All one can reasonably hope for is that the proper direction has been pointed out and that this book will find a receptive audience who will concur with what, ideally, has been shown in retrospect to be plainly obvious.

Among all disciplines, I've noticed, there's a tendency in philosophy to prize traditionalism over innovation, to favor the notion that its problems and questions are timeless and immutable in their as-given forms. Perhaps the motivational payoff for such disciplinary conservatism is this feeling that one is engaged with the eternal questions, the fundamental mysteries which, if we meddled with the essential formulations in which they've been cast, would cheapen and lessen the allure of their inexhaustible profundity. So there is this implicit attitude that no updates or revisions can be made about what the foundational concerns of philosophy are, but rather only new arguments can be posed about the same problems. One is tempted to wonder if the same circles have been retraced for centuries not because these deep questions are essentially insoluble or that philosophy itself is ill-equipped to face the challenge. One wonders instead if the lack of incontrovertible progress is attributable to the way the problems have been framed in their set in stone formulations. To that end, it turns out that epistemology is missing critical pieces of the puzzle, and until fairly recently nobody noticed these pieces were missing or counted them as relevant.

Epistemologists unintentionally spite Socrates, they know they know something, and that something is knowledge . Except epistemology is a more expansive project than "simply" defining knowledge, or purporting to name what conditions justify claiming the possession of it, or describing how it aligns mind and world together in correspondence with the truth, and the rest that follows. All this is good and we should keep going at it. What's missing is a description of the proper setting in which to situate this study of positive knowledge, the context of its full *etiology*, its reason for existing, which is achieved by adding an understanding of ignorance and inquiry into the mix. With deep irony, epistemology is ignorant of ignorance; ignorant of its own self-

limitation as it is presently conceived. This ignorance blindspot paradoxically consists of both the solution and the problem, the cause and the effect, the target and the source, of exactly what it needs to understand to progress beyond its ancient preconceptions.

Appendiix I: Responses to Anticipated Criticisms

1. Your book should've been a much shorter. Ignorance actually is just the negation or absence of knowledge.

Probably the most predictable critique is the same reason why epistemology has been so dismissive of ignorance as a subject in the first place: it really is just the negation of knowledge. Or it is the absence of knowledge, which is not the same as negation. So efforts to read too deeply into ignorance are frivolous. Much better to understand knowledge and then simply apply the negation operator to all those results to understand ignorance, or one supposes, to imagine what would be to not know it, picture a sort of contentless and meaningless mental blankness. Negation is not erasure or nonexistence, it asserts the inverse of a truth value : if P is true then ¬P is false and vice versa. So on the face of it, there are two distinct concepts here, "it is false that it is known" and a lack of awareness. There are many ways to argue that ignorance is more than the negation of knowledge, but using epistemic logic is the only way to *prove* it. This can be achieved by two separate proofs. The first demonstrates that ignorance considered as uncertainty is not the same as ignorance considered as the non-knowledge. Unfortunately this argument must be less rigorously demonstrated, because epistemic logic as far as I know has no means of representing degrees of belief or degrees of certainty. The second shows that ignorance can be identified with a false belief and therefore a propositional attitude that is false awareness and not a lack of awareness.

a. Premise: Ignorance is more than the negation of knowledge.

Proof: knowledge, by convention, is always true. The standard truth axiom

in epistemic logic Kbϕ → ϕ, states that if b knows ϕ then ϕ is true. Certainly , it is also true however that :

¬Kbϕ ↛ □ ¬ϕ

if b does not know ϕ that does not necessarily imply ϕ is false. The truth of ϕ is not conditioned on b knowing it at all. Conversely, whether one believes ϕ is true is independent of whether ϕ is true. Likewise whether one believes ϕ is false is independent of whether ϕ is false . This might be expressed:

Bbϕ ↛ ϕ ∧ Bb¬ϕ ↛ ¬ϕ

One might call this the *arbitrariness axiom* of belief. It is what differentiates it from the truth axiom of knowledge which again states if one knows ϕ it must be true. Belief therefore is different from knowledge in that its degrees of freedom are less strict. From this axiom it can be seen that one can have a true belief and still not have knowledge, because the truth axiom of knowledge does not apply to belief. This brings us to Gettier's infamous problems ([https://web.archive.org/web/20170324194044/http://philosophyfaculty.ucsd.edu/faculty/rarneson/Courses/gettierphilreading.pdf]).

While it is necessary—at least by convention—that knowledge entails true belief, a true belief is not sufficient to entail knowledge. So true belief *can* be not-knowledge, one can incidentally believe the truth without knowing one knows it , failing the so called positive introspection axiom. But this also means that one can believe one knows or know one knows and still be right in both cases. Does this entail that believing one knows dilutes the knowledge? What it does is that it injects a degree of uncertainty into it. If one believed one knew and then was shown to be right, one knew.

This is a point at which it is more helpful to step out of the strict rules and look at a real world example. Because epistemic logic has no concept of causality —I argue it can't unless it expands its repetoire to accomodate for ignorance and inquiry operations—it is impossible to reason accurately about knowledge without assuming unrealistically perfect conditions for reasoning, so called logical omniscience.

If one believes one knows who will win a ballgame, and they win the ballgame, one can retroactively and with good grounds amend one's prior belief

and justify it as knowledge. If one knows who will win, the result must strictly be that you were right from the beginning , impossible for you to wrong, thanks to your godlike powers of prophecy. Depending on the consequences one can justifiably promote prior belief to knowledge. Before the results of the game are confirmed, the only thing that differentiates belief or knowledge in this case is degrees of certainty—how much one believes in the proposition. Certainty must arise from somewhere, and it is clear that one cannot know without certainty. Only beliefs may be uncertain.

For the believer, once the results are confirmed to be consistent with their belief, one's level of certainty in it is upgraded to the status of knowledge. The fact of who won is absolutely certain. Can one truly know what the result of the game will be with perfect, rational, and totally justified certainty? The only true justification of either knowledge or belief is that the result confirms it.

If one person believes they will win and another one knows they will and they win, they are equally justified in claiming they knew it. However we somehow still maintain the prejudice that belief is epistemically weaker than knowledge, but this is only because of how we view them as they play out in time. The truth axiom may apply in the idealized setting, but if it does not apply in the real world it cannot be said to be an actual codification of a law of knowledge.

There is no reason to suppose that justification must come *before* the confirmation of the fact, in which case the only difference between knowledge and true belief is that justification comes before or after.

Only if one is proven wrong by the result can belief then be seen as inferior to knowledge. Suppose that one of them believes they know who will win and the other knows they believe who will win and the team loses. Which of these is the more epistemically weaker position? The answer is neither. Both afford themselves the possibility of being wrong. As soon as belief touches knowledge, knowledge seems to lose all grounds for validity, This of course, is an absurdity, since knowledge entails true belief.

The difference ultimately is only in a subjective sense of certainty, the believer is not prepared to know until they gain confirmation. At which point however

they can't be said to not know; they have only conditioned their knowledge on the outcome, rather than clairvoyance or whatever means the knower uses.

In the former case involving belief of belief, however, no negation signs are introduced, only two affirmative assertions are made. Yet *the conjunction of two true propositions* is always true and cannot be false. At no point is a false belief introduced that would immediately disqualify it from being knowledge once confirmation is reached. If ignorance is strictly not knowing then, the negation of knowledge, at some point in this formula a negation sign must result, a belief must be falsified. Whatever differentiates a true belief from knowledge then is not justification, the only true justification is proof of outcome. What is ignorant about the unknowledgeable true belief is not that it it contradicts anything that the knowledge predicts. It is that the belief does not *rule out possibilities* In other words, information entropy.

Uncertainty then is different from not-knowledge, in that to be uncertain is to essentially allow that there are possibilities ,*to know one believes* that there are more than one outcomes. whereas to not know means to assert what is false. To know one has an uncertain belief is to be certain one is uncertain. Ignorance then is openness to possibility without specification of any outcome rather than the assertion of an outcome that is false.

And a true belief can be uncertain in the sense that it knows that either an event will happen or it will not happen. Ignorance as uncertainty *carries information* in terms of alternatives, in the manner described by Shannon entropy.

knowledge then, comprises a subset of true beliefs distinguished only by certainty. Establishing grounds for such certainty outside of the closed worlds of deductive systems, however, is not demonstrable by any universal rule. So real world knowledge is always more like true belief with a high grade of certainty rather than strict knowledge.

This brings us to the ultimate point. Is the belief that one is uncertain any different than the knowledge one is uncertain? Suppose that the proposition ϕ states that "it is uncertain that." Then it can be said:

$Bb\phi \equiv Kb\phi$

Believing that it is uncertain is logically equivalent to knowing it is uncertain, Because ϕ is no less uncertain whether b believes it is or whether b knows it is, according to the axiom of arbitrariness. If ϕ and B believes ϕ is false, it is *no less* uncertain. The belief cannot avoid uncertainty in either case. If b knows it is uncertain then it is just as uncertain as if they believed it was false whether it was uncertain.

Now if b knows it is uncertain, then it is just as uncertain as if they believed it was uncertain or disbelieved it was uncertain or knew it was uncertain:

$\phi \rightarrow B_b\phi \wedge \neg B_b\phi \wedge K_b\phi$

What this means is that all three conditions, true belief, knowledge, and false belief, equally express the same fact of uncertainty with equal truth. But this would mean that a false belief is just as true as knowledge. The only way to avoid a logical impossibility then is to assert that ignorance is represented by all of these conclusions.

The only way to disprove this is to assert that it is impossible to know that a given state of affairs is uncertain and knowledge strictly implies certainty about its proposition. There can be certain knowledge of uncertainty so nothing is accomplished by this counterargument. It is only in the much criticized and frankly fictionalized world of logical omniscience that it is impossible to be uncertain. Therefore it must be true that ignorance is at least both the negation of knowledge and uncertainty.

The poverty of epistemic logic is that it is incapable of comprehending probability, because it is incapable of comprehending ignorance. Uncertainty expresses an immutable modal quality which must be connected to notions of entropy.

knowledge and ignorance-as -uncertainty are therefore differentiated by degree rather than completely in terms of how many possibilities they rule out. Ignorance as uncertainty then is distinct from ignorance as lack of knowledge.

b: Ignorance is the absence of knowledge

*Argument *

Ignorance is much easier to define in epistemic logic when considering false beliefs.

An ignorant belief has a specific form. To express the form of ignorance one could write :

ϕ, Bb$\neg\phi \rightarrow \neg$Kb$\phi$

ϕ is true, b believes that ϕ is false, therefore b does not know ϕ. This would imply:

ϕ , Bb$\neg\phi \leftrightarrow \neg$Kb$\phi \vee \neg\phi$, Bb$\phi \leftrightarrow \neg$Kb$\phi$

And inversely

Bb$\phi \leftrightarrow$ Kbϕ

Which essentially says you can only not know if you believe what is true is false (or vice versa) and you can only know if you believe a true proposition is true. This is nothing outrageous, however, it effectively shows that ignorance is not simply the negation of knowledge but is also definable as false belief, that it is dependent on belief. Belief is not the mere negation or absence of knowledge, it holds a propositional attitude, a content. This can be more clearly illustrated if we look at it in terms of the implicational relations of the premises:

$\beta \oplus \phi$, Bb$\beta \rightarrow \neg$Kbϕ

Either β or ϕ but not both is true. B believes β, therefore B does not know ϕ. However this shows that a false belief can be affirmed without negating ϕ. Meaning, one can believe falsely without believing the opposite.

Believing wrongly can mean believing something other than Kϕ, it does not mean believing $\neg$Kϕ. This is obvious from the fact that someone who believes something false does not believe the opposite of what they would believe if they had knowledge, such that if they merely added "not" to what they believed they would know the truth, by applying a so called symmetrical accessibility relation. One can falsely believe that all humans are vegetarians, but one would be equally wrong if one assumed that no humans are vegetarians. So reversing a false belief does not arrive at a true belief. That could, in principle, happen, but more likely their false belief is one among many possible false beliefs while there is strictly speaking only one true belief. Hintikka (1962) was onto this point with his notion of epistemic alternatives.

Various logical structures of ignorance can be defined.

$\beta \oplus \phi$, Bm$\beta \wedge$ Kmϕ

This says that either β or ϕ is true, and m believes β and knows ϕ. This essentially is the logical form of *confusion*. However since a belief is under no obligation to be true, there is nothing invalid about this statement as would be the case if it was expressed as Kmβ ∧ Kmϕ. The ignorance here then is in the failure to relate the premises correctly, to allow falsehoods to going along with truths. Something like the *complex ignorance* I've mentioned.

ϕ, Bm Km¬ϕ → ¬Kmϕ ∧ ¬Km¬Kmϕ

Might be taken for a model of *delusion*. It states ϕ is true, m believes they know ϕ is false therefore , m does not know ϕ and does not know they do not know it. Belief in a falsehood is not an absence or negation of a truth, is is the assertion of a falsehood. While not knowing is certainly involved here, as it always is with ignorance either directly or indirectly, it plays more of a regulatory function rather than constituting the substance of the misunderstanding on display.

These examples show that when you mix beliefs and knowledge it becomes harder to reduce the ignorance to not knowing.

The one axiom in epistemic logic that comes closest to touching ignorance is the negative introspection axiom:

¬Kbϕ → Kb¬Kbϕ:

if b does not know ϕ, b knows that it does not know ϕ. Reading this as an English sentence defies all common sense. Considering the idea that ignorance is an absence of awareness of a proposition it is quite clear that this is not a necessary inference. This raises the specter of logical omniscience that is often leveled against epistemic logic, a problem I would argue only arises because it lacks an ignorance operator. The axiom of negative introspection has been criticized as assuming an unrealistic closure property of epistemic agents, that even if b does know the proposition it always inspects it through the use of the K modal operator and can therefore determine that it does not know it. (Hendricks, Vincent F., 2005, Mainstream and Formal Epistemology, Cambridge: Cambridge University Press. doi:10.1017/CBO9780511616150)

The best definition of ignorance as absence of knowledge would then be the negation of the negative introspection axiom. To spare myself from writing it out,

this would result to: b does not know that b does not know ϕ.

The result is a very inelegant way to express a very simple concept. All of these pitfalls and loopholes invariably stem from an incomprehension or devaluation of the logic of ignorance. Logicians working on epistemic logic have tended to favor conditions of ideal inferential salience that basically force all negative concepts out of the picture. Solutions have been proposed to avert these misrepresentations, most notably for our purposes the introduction of the awareness operator (Fagin and Halpern 1987[https://www.cs.cornell.edu/home/halpern/papers/awareness.pdf]). The awareness operator differentiates between implicit and explicit knowledge depending on whether the agent is aware they know or not. Awareness, however, merely pushes the problem back further, because awareness is a veiled form of knowing, except that it is purely introspective in its us as an operator. It would seem then, that the true logical expression of ignorance would be not knowing (semantically) *and* not being aware (consciously), folding them together in a single operator I. Aversion to this idea probably stems from the belief that only affirmative terms should be counted as primitives because they they are clearer, and that an operator should not express two meanings. For one thing the ignorance operator is an affirmation, it *affirms* a state of ignorance, and it saves one from having to overuse negation symbols, which make things harder to process. Secondly, it is only an illusion of language that the ignorance operator would be considered to have two meanings, it designates the single meaning of ignoration, total absence of knowledge in all its traces. For the sake of syntactic economy alone to spare us from having to clutter up pages with negation signs, an ignorance operator is worth introducing. On a more advanced semantic level, it would allow for a more realistic logic and it would banish all traces of logical omniscience.

And that would be right. Had I the inclination I would write out all the formulas for the different structures of ignorance, like the confusion and delusion examples listed above, and pin them up like butterflies on corkboard. Now is not the time however. The point is that documenting these structures is not a silly pastime, is no more or no less valuable than documenting knowledge structures, and doing so balances both sides of the epistemological universe. Keeping

epistemic logic pure of ignorance is all well in good so long as we at least extend it to a new branch that can accomodate ignorance.

Much more would have to be said to develop these notions, but this is not the place for it.

Moving beyond formal considerations to a more general restatement of it. Taking ignorance as the negation of knowledge does not explain how ignorance affect beliefs. An unwitting believer in falsehood has wrong reasons for believing it, it is not just that they lack the right reasons. They still have a reason, they are not believing a falsehood just because they don't have a reason not to believe it. A delusion is more than simply the lack of insight into the irrationality of a belief. Such ignorance has a power to misrepresent reality. Further, the belief that ignorance is just the negation of knowledge fails to explain why there are so many subtypes of ignorance that can be convincingly identified and organized into a taxonomy. Is this an idle practice? Assembling a taxonomic tree of detailed and distinguished yet imaginary entities, just like the medieval scholastics did placing the varieties of angels into a hierarchy? Or does the capacity to organize well defined subtypes of ignorance point to actual identifications? While not all taxonomic categories of abstract types are equally well founded, the ability to distinguish types of ignorance demonstrates they cannot all be not-knowledge. Given that I have provided good reasons for each of taxonomic distinctions, I stand by the claim.

2. The claim that ignorance and uncertainty is somehow objective commits the mind-projection fallacy: the fallacy which confuses how the mind operates with how the world is.

Before proceeding, I would like to clear the record and say that I'm not arguing that ignorance literally exists outside our heads, although neither am I ruling it out. Only that it is insightful to think about it this way, to view it as a kind of cognitive dark matter that is simultaneously part and not part of the mind.

This critique somewhat ironically originates with E.T Jaynes, the physicist

whose ideas are one of the pillars supporting this book's whole argument. It also happens to be the most compelling. According to Jaynes,

"In studying probability theory, it was vaguely troubling to see reference to "gaussian random variables", or "stochastic processes", or "stationary time series", or "disorder", as if the property of being gaussian, random, stochastic, stationary, or disorderly is a real property, like the property of possessing mass or length, existing in Nature. Indeed, some seek to develop statistical tests to determine the presence of these properties in their data… (p 3)"

(Ninth Annual Workshop on Maximum Entropy and Bayesian Methods. pp. 1–
16. CiteSeerX 10.1.1.63.5075. doi:10.1007/978-94-009-0683-9_1. ISBN 978-94-010-6792-8. Retrieved 2011-05-19.)_

Like Einstein, Jaynes thought that "God does not play dice" and was skeptical of the assertions, many of them propounded by the original pioneers of quantum physics, that quantum behaviors are by nature indeterministic. This attribution of indeterminacy, thought Jaynes, mistakes the mind's inability to comprehend an underlying, hidden order for the way the world works. These probability models are tools of thought, or more like coping strategies for what we can't fully understand.

At first glance this criticism, coming especially from the same person who provided the conceptual framework that informs a major part of this book and who was expert in both the mathematics of probability and physics, appears fatal. How can ignorance be seen as a *something*, existing in any way, shape or form outside our thoughts and their relationships to things? How can it be a kind of substance, suspended in a manifold alongside the tangled ball of yarn of the bits and pieces of knowledge we possess? Even students of agniology would shy away from this act of gross reification.

I'd respond by saying that the premise that the mind and the world are somehow ontologically alienated and demanding of entirely separate descriptive modes is a profound error endemic to much of western thought, traceable perhaps to Descartes' mind-body dualism. What is true of the mind can also be true of the world and vice versa. The way minds think is a consequence of the

particular nature of this world. If the mind can discern the truth at all, then it must share some property with the world at large that enables it to tell what is real. No matter how warped and riddled with biases this mirror may be, it still reflects. The subject-object distinction is misleading as a basis for grounding all philosophical discourse. This is, after Wheeler, a participatory universe.

Not that I propose some far-out panpsychist take on things, even though if the universe is essentially mental in nature this would provide a perfect justification for the objectivity of ignorance, as there is no mind in the physical universe that is free of information incompleteness as per Wolpert. (Although it's worth noting that panpsychism does have respectable proponents such as Chalmers and Goff. I am not discarding it out of hand.) Nor am I casting my lot with solipsism or some trippy Schopenhauerian sentiment that the world is a dream or "my idea", and that the object is an illusion and so my knowledge and ignorance are inseparably interwoven with perceived external reality. I propose rather that in the absence of a mind, the universe itself is undefined, maximally ignorant because no part of it such as us has awareness. Thus, ignorance is to some extent separable from the existence of an individual mind, and therefore must be a universal objective phenomenon. With no one around to see it, the cosmos is no different from total oblivion and total ignorance is no different from it. The ultimate ignorance is the complete absence of mind from the universe, the big sleep. It is pure thoughtlessness. My counterargument is a version of the anthropomorphic principle or Kantianism. To the extent that anything has an objective property such as mass or length, it is an observation of a mind. After all what is a length without comparing it to another length, and the same for mass? Longer or shorter are relative terms, and it is the mind which specifies these relations. It is ultimately the mind which proportions objects. Stochastic processes or random variables are themselves expressions of unknowing, and so it is an odd detail indeed that concepts considered to be part of the repetoire of scientific knowledge are themselves the stuff of ignorance. If these concepts merely describe our ignorance as to the objective state of affairs, they describe nothing about the state of affairs. So why keep them around in the science? A deterministic world ought to permit perfectly deterministic mathematical models

without fuss. In the very next paragraph Jaynes says that he was able to solve "very real problems" by applying these mathematical constructs. So what gives? For an idea to be applicable, for it to achieve results, mustn't it reflect some underlying truth about the world in the same way that applying the principles of engineering correctly results in a sturdy bridge? Why is probability useful if it is only about mental figments?

Jaynes sums up the mind projection fallacy with this elementary formula:

(A) (My own imagination) —> (Real property of Nature)

(B) (My own ignorance) -> (Nature is indeterminate)

Inference (B) is at the root of the matter. "Form (B) arose out of quantum theory; instead of covering up our ignorance with fanciful assumptions about reality, one accepts that ignorance but attributes it to Nature." (p.3) Yes, as one should. Nature is not some categorically distinct and separate entity from the minds embedded within it.

While it proves nothing it's worth pointing out that Jaynes is among a minority of physicists who think this way. Quantum mechanical phenomena are simply immeasurable without resorting to these probabilistic formulations, and for a physicist to quantify is to take the measure of the facts, to obtain truth. At very least, if physics is the quantification of nature, then there is no way to do so without resorting to these indeterministic equations unless you posit an even more lavish assumption such as an unobservable multiverse. That happens to be because multiverse theory explains away the probabilistic behavior of the wave function by claiming that rather than collapsing into one observed state from the superposition of all possible states, the particle takes all of them at once and splits into a different universe for each state. Is that any less of an extraordinary assumption? (Vaidman, L. Quantum theory and determinism. Quantum Stud.: Math. Found. 1, 5–38 (2014). https://doi.org/10.1007/s40509-014-0008-4)

So what is it then, an inexplicable multiverse or self-subsistent indeterminacy? The decision between them is itself indeterminate, because it can never truly be observed that the multiverse exists because it is theoretically impossible for different universes to communicate. Neither can a hidden determinate causality be ascribed to the indeterminate probability states. Additionally any

indeterminate process can be observed in retrospect to have displayed a definite causal history, even though as it unfolded prior to measurement this information was impossible to discern.

There are more questions. Why is information theory so robust if its object of measurement is uncertainty? How is it possible to measure a vague mental state, which is not only vague in what it pictures but *consists* of vagueness, in terms of bits? Why is modern physics so informationalized if it is based on a theory that quantifies unreality, a tissue of mental fictions? Is entropy, as the measure of the disorder of a physical system, a false concept? Why then is the science of thermodynamics concerning itself with such games? Jaynes appears to conflate indeterminacy with the absence of causation, which is not a necessary inference. Something can indeed be *caused* to behave randomly and unpredictably, as with a broken electronic device that glitches unexpectedly. Causality may be determinate but phenomena can still be indeterminate. One can hook up a machine to a random number generator and have a red light go on every time a multiple of 3 comes up. While an RNG is not purely random, it approximates it, and it still is what causes the red light to go on. Simply that randomness has a cause does not negate the reality of that randomness or imply that it has a hidden order we cannot detect. And in this case the burden of proof is on determinacy. The indeterminate equations speak for themselves. They fit the observations.

The problem with equating causality with determinacy is that the source of causation is obscured. Causality is itself apparently uncaused. It exists for whatever indeterminate arbitrary reason for why causality happens to exist. It originates that is to say from an unknowable unknown (or from a prime mover such as God, if that's your taste). Alternatively as per Hume, causality is not what we think it is and all we really truly observe is the temporal and spatial contiguity of objects in certain reliable frequencies of repetition and circumstance. Causality in itself is not observable, it is loosely inferred.

The counterargument then, can be summed up by saying that to the extent that we can only think accurately at all about a given phenomenon is in terms of probability, then that is just as good as saying that the probability is a property of the phenomenon. To the extent that Kant was on the right track, it is impossible

to know how things are in themselves and all reasoning involves mind projection. *The impossibility of knowledge of the Thing in Itself is therefore a property attributable to it* to the degree that an intrinsic property can be attributable to anything. If it is an inevitable result that all possible minds have the same reaction to something, what they are responding to is a fact. Thus the unknowability of the thing-in-itself is a fact over and above any given mind.

Because ignorance is inescapable and originative , and since no matter where we turn we find it waiting, it has intrinsic reality even if it is unlike anything and is not literally floating around like visible balloons. If the incompleteness of our knowledge and the available information is an objective constant which presupposes the flawed nature of inductive reasoning and the laws of physics limiting information access, and if reasoning is the means by which we differentiate the objective than the subjective, the real and the merely believed or thought, then we are either bound to the conclusion that all knowledge is just arbitrary opinion or that ignorance is a fact of nature. Our very means of determining the truth cannot escape ignorance.

When all is said and done it a perfectly deterministic universe is consistent with my arguments so long its causal mechanisms inevitably lead to obfuscation. If such a causal system's inscrutable processes reliably cause ignorance by screening its operations through a veil of apparent disorder and irregularity then this is just as good as saying that nature is determined to reliably cause ignorance. If the appearance of probability is caused by overwhelming deterministic processes, then those appearances are still not a figment of the mind but a natural way to quantify and organize unintelligible hidden variables.

3. If ignorance is objective and exists "outside the head" as part of the external world why not knowledge? What's the difference?

knowledge can be, and sometimes is external, although not nearly in such abundance. It can exist outside the mind, or at least artifacts that represent exteriorized knowledge can. There's just a lot less objectified knowledge considering the vastness of the mostly inanimate universe, but what is there is

much more concentrated and explicit. It must be explicit because we are who create it. knowledge does not just show up the way ignorance does, primordial and inertial since forever. A microchip, a cathedral, a spaceship, a bank, an AI program, an x-ray scanner, these are all objectified knowledge that has been "encoded" into physical materials or embodied in groups of coordinated brains and human bodies. Any result of the impulse of intelligence to organize, sort, order, construct, describe, document or implement that leaves a trace outside the body is knowledge embedded into the world. When behaving intelligently we straighten out the world, impose a workable order on primordial anarchy. This is not a new idea. Indeed, the idea that cognition can be distributed and extended across brains, bodies, groups, and tools is a fairly well established strain of cognitive scientific thought (Cognition in the Wild, Edwin Hutchins). People excrete knowledge onto their environments like bees do honey filling up the honeycombs that are libraries. At which point this knowledge enters a resting, stored state like an executable program in memory storage waiting for the go-ahead to fire up. Efforts to reduce information entropy lead to concentrated ordered states within our living environments and work settings which we use to guide ourselves through the fog of prevailing uncertainty beyond our island bubbles of relative predictability.

A messy office desk atop which it's difficult to find anything is a configuration of matter that is ignorance (namely confusion) inducing. The irregularity of the stuff on the desk causes ignorance. Inversely, a well organized and tidy desk has become a system, a temporary one granted, one which will tend to dissipate into gradual chaos if it is not kept up with, but a system nonetheless. Certain states of the world, like a book or presentation perhaps, do constitute knowledge inducing stimuli, however these signals predominantly come from other knowledge agents. There is more ignorance causing physical states in the world because the organized desk is the inevitable exception. Effort has to be put into tidying a desk but the natural tendency is for things to get trashy, to accumulate "kipple", entropic waste buildup. Finally, the distribution of particles of the universe is nonuniform. Because of this skew, we cannot apply efficient symmetric bijective inferences to it, we cannot say that what is true of one half is

true of the other, and so on for any arbitrary quadrant of spacetime. At least, that is, not in the details.

4. Talk of including ignorance and inquiry into epistemology bloats it. Epistemology is about knowledge on its own terms, plain and simple. What you seek to create is not epistemology proper but a philosophical Frankenstein .

Admittedly this is more semantic hairsplitting than a credible refutation of my proposals, but it does carry significance and cannot be dismissed. How we package ideas together or keep them separate influences how we think about them. It can help emphasize their relationships or it can obscure them.

This criticism is simply that epistemology is about the inner nature of knowledge in and of itself, and how knowledge states link to other knowledge states through properties of closure and inference. It is confusing to stuff it with additional categories of ideas that do not share these same class membership criteria. Keep things neat. File the concepts of inquiry and ignorance under different headers and leave epistemology undisturbed.

No, because If the argument that epistemology is deficient so long as it only references positive knowledge is true, there is every reason to include with it the ideas that would correct this flaw even if they operate by different principles and it would complicate the picture. Sometimes things need to be complicated if that means a body of explanation can then deal with greater nuance and complexity because its universe of discourse has been expanded to accomodate more reality.

Epistemic logic would benefit from an expanded notion that formalizes ignorance and knowledge-ignorance relations more thoroughly. To my knowledge the only rule of epistemic logic that enshrines the recognition of such a relation is the negative introspection axiom, which essentially formalizes the statement “M knows that they do not know that they do not know.” This alone is not much to chew on. It simply ties off a loophole that would otherwise be the inroad into ignorance. Hintikka was on the right track by attempting to extend epistemic logic to encompass the questioning process, what is also needed is an

“agniologic logic”, a logic of ignorance.

The axiom of negative introspection appears to hold a structural function in epistemic logic, plugging up a hole so that no more consideration of ignorance need happen. Ignorance of course is much broader than the one subtype of simple not knowing (\~K), but if this is epistemic logic’s only definition of it that would explain why it has such an impoverished conception. It terminates at the simple. It stops at negation, which is secondary to assertion. One suspects that the disinterest in ignorance is wound up with the widespread belief that negative statements are less informative than affirmative statements. If that condition doesn’t hold and ignorance in fact has a more detailed structure than just \~K, much more is to be gained by including it.

knowledge can no more exist or be understood without ignorance than hot can be conceived without cold. As for the theory of inquiry, knowledge and ignorance cannot coexist without specifying the means by which they are connected, the means which explains how that knowledge came to be. Specifically a *causal* logic of knowledge and a theory of its acquisition must reference ignorance.

Lastly, putting the study of ignorance into its own category, agniology, mentally biases us to think of it as unrelated to the study of knowledge, risking a reversal of the same incompleteness problem in epistemology except the theory of ignorance might be knowledge-blind. And from a “marketing” perspective attaching ignorance theory to the big name brand of epistemology is a good way to raise awareness and popularize it rather than to tuck it away in a less prominent, perhaps even fringe field of philosophy.

5. Even if it makes sense to include ignorance and inquiry from a categorical point of view, adding more concepts we don’t completely understand to a domain we don’t completely understand can only increase the overall lack of understanding in that domain.

This point expands on the previous. It is true that the theories of inquiry and ignorance are relatively nascent, and that merging them with the theory of

knowledge just gives us more work to do. Doing so introduces a host of new problems. But this misses the point. Unequivocally: *There will never be a complete understanding of knowledge without an adequate understanding of these other concepts.* We have no choice but to bite the bullet and figure out how they go together. Harmonizing these three categories recasts old problems in a new light. At best, it removes fruitless problems which arise solely from a misunderstanding of the matter and adds problems which are productive to solve.

To combine things that share special affinities, even though they are not the same can trigger a whole shift in perpective. Once the conceptual space of epistemology is widened in scope, new connections can be formed which were inconceivable before.

Any way you slice it, knowledge is still central to epistemology. There can be no ignorance-first epistemology. All of this orbits knowledge. You still haven't defined knowledge. We still don't know what knowledge is with certitude. So it's pointless to study ignorance unless we can obtain knowledge of knowledge. There's a hole at the core of your study.

The knowledge of knowledge demands knowledge of ignorance as a prerequisite. From a causal point of view and from the standpoint of its origins, knowledge cannot come before ignorance. Everyone starts out ignorant, and ignorance always precedes knowledge. All discursive reasoning "targets" ignorance, or else there would be no point too thought. All problem solving is a reaction to ignorance, or else there would be no problem; we'd already know what to do. Ignorance is undeniably the primordial and precursor condition from which knowledge emerges. Granted to understand ignorance we must come to know it and subordinate it to knowledge, but this knowledge is but a mere pocket within a vast expanse of ignorance within which it gradually extends its tendrils and grows its network of associations. It bears mentioning also the definitional problem of knowledge has not been solved within epistemology as it

presently stands. So nothing is jeopardized by broadening the subject matter of the field. If anything the addition of fresh ideas might pave the way to a solution.

Contrarily, the definition of ignorance is much more easily expressed. It is easier to claim with certainty how or that we don't know rather than how or that we do know. There is never any doubt that we are doubting when we are doing so. Doubt itself cannot be doubted, as that simply leads to more doubt. A state of knowing is more brittle, we can come to doubt it if reasons arise. While doubt may subside into a state of belief or knowledge through the completion of a process of reasoning, it may again dissolve back into doubt if evidence that cannot be reconciled with the belief shows up to challenge it. When it comes to identifying our ignorance there is no "maybe." When we are in for a rude awakening it's usually because what we thought we knew was wrong. It's indeed a rare and pleasant surprise to think we don't know and then to suddenly behave with miraculous savant powers. Though this is possible it is less probable than the alternative case. Ignorance in itself is uniformly contentless, whereas knowledge has a complicated representational content that is at least as complex as the brain that produces it. (Even given that there is such a thing as complex ignorance, this type of ignorance is tacitly embedded in representations as a corruption of stateful beliefs. The complexity comes from belief system rather than from the ignorance misguiding it behind the scenes.) Though the definition of ignorance is somewhat dependent on the definition of knowledge, the reverse is also true. So a solid definition of ignorance helps to balance the definition of knowledge.

I did not dedicate a section of the book to traditional epistemology because there is a bountiful catalogue of books, treatise, articles, academic journals and other materials on the subject going back more than two thousand years. The literature on ignorance is comparatively much scarcer. I doubt that I would say anything that hasn't been heard before if I did treat knowledge directly, at least not without reference to ignorance, and in any case, it would not be conducive to my stated aims to do so.

7. You underestimate the power and scope of knowledge. People have figured out quite a bit over the ages and it's only speeding up as new knowledge indeed grows off old knowledge. We collectively and at times individually know quite a bit and ignorance does not always dominate our societies and lives. It's both pessimistic and false to suppose that we may never know the ultimate secrets of nature or that we will always be bounded by some level of ignorance.

It's important here to stress the idea of quality over quantity. knowledge is indeed powerful, but it is rare and obtained with difficulty. There is a kind of foreshortening optical illusion effect where, since we cannot see how much we don't know, we overestimate how much we do know. What is common and requires no effort is necessarily what's most frequent. It is indeed an absolute truth that it is impossible to know everything and that most facts will never be known by any one of us. The impossibility of omniscience therefore strictly entails that there will always be more ignorance than knowledge. Everyone is utterly oblivious about the vast majority of what is the case. Except most of those facts are junk, useless trivia or redundancies without any meaningful context and only some of them are important to us. As I made explicit elsewhere, the best knowledge is pressed into the shape of universal theories which generalize about the invariant features of a diverse array of related phenomena. This helps us to think better rather than waste time thinking more, to evaluate a member in terms of a category rather than a category in terms of each member. While this is not a perfect method, when done right it leads to the most vigorous kinds of knowledge. An ounce of knowledge is worth a pound of ignorance; the quality of the former far outpaces the quantity of the latter.

I call for realism, not pessimism or optimism. It's simply a brute fact that more truths are unknown and some truths are unknowable from our standpoint in space and time. Some of them may matter more or less than others, but there is quite a lot either way. By the same token it is itself unknowable whether science has already discovered all the deep fundamental theories of nature and merely has to work out the details. It is impossible to scope out what might be left. They may yet be "nomological danglers" inconsistencies and loose ends in the

overlapping superstructure of natural law. knowledge about the extent of our knowledge—what epistemic logic codifies in the positive introspection axiom—is constrained by conditions of undecidable ignorance. The ratio between the known and the remaining unknown is an undefinable quantity. The best we might hope for is *consilience* as biologist E.O Wilson calls it, the knowledge that all our knowledge is aligned and non-contradictory. Each branch of science is consistent with the others, but internal coherence alone does not prove that more potential knowledge is still outside that scope.

That said there is a reason a "know it all" is a term of derision. Overconfidence lulls us in self-imposed certainties that may have no bearing on the facts. The worst mistakes, the gravest miscalculations, the horrors of violent ideologies above all are enabled by unshakable self-certainty and a drastically inaccurate overestimation of knowledge.

A literal know-it-all, moreover, is a logical and physical impossibility. Such a being would need infinite knowledge, including an infinite knowledge of everything that is possible or impossible or necessary in all possible worlds, and infinite knowledge that it knows each of these propositions, and this is impossible for any finite being. If the universe is infinite, the know-it-all would have to be an infinite being with infinite knowledge that knows itself infinitely, and supposing it to have some mechanistic or algorithmic means of achieving this would run up against computational and therefore logical limits. If the universe is finite, then any finite being must know fewer facts than there are about that world, as Wolpert's work formally rules out. Unless were are happy to accept a Spinozist definition of God (in Ethics, he defines God as, "an infinite substance consisting of infinite attributes each of which expresses an infinite essence") as the one example of omniscience and omnipresence, then we can rule this out on logical and physical grounds. Even so, the existence of God is neither a logical necessity and if God were to exist, God is not to be found in any finite physical system of the material universe. That said, a finite being has finite needs. We do not need absolute and complete knowledge to have adequate knowledge to achieve our ends. And while it is unjustifiable to point to any given unsolved problem and say "it can't be known" it can be said that literally everything can't

be known. Neither can it be asserted a priori that all truths must be knowable, unless you already know all truths (an impossibility), so the truth of this claim is equally as undecidable as the claim that some truths are unknowable. *This* however, is a good candidate for the existence of at least one unknowable fact!

It is not a weakness or sign of cowardice to err on the side of caution and recognize the prevalence of ignorance both within ourselves and at large but to instead a profitable exercise that optimizes learning and growth.

8. Because ignorance is by definition inscrutable and vague, even as a concept rather than as an instance of it, you can attribute any properties to it arbitrarily. You reify ignorance, attributing excessive causal powers to it and nearly suggesting that it has more complexity than knowledge. How is it possible to argue for *or* against these assertions when there is nothing tangible to which we can refer? Given the inscrutability of ignorance, you can suppose any number of things about it none of which can be ruled out. Nothing you say about it is falsifiable. Occam's razor would favor the most simplest explanation in these circumstances (ignorance is not knowing)

The accusation that this entire treatment is just seeing castles in the clouds cannot be entirely dismissed, because the central argument that ignorance must be understood before knowledge, has structure, or is even ontologically subsistent reality, like any bold claim deserves skepticism. There is also validity to the point that loose thinking can impute anything to abstractions without no well-formed inference rules to differentiate between justified and unjustified claims.

The seeming paradox of understanding ignorance is a category mistake. "*How can I know that which I do not know? Must not any attempt to understand ignorance alter it, and any successful attempt destroy it?* . asks DeNicola. No, because this is a confusion "between understanding the *concept* of ignorance and understanding *that which we do not understand* (DeNicola 2018, p 19). The concept of ignorance is intelligible.

It is trivially true that ignorance has no real physical or sensible properties,

nothing to be put on a scale and weighed. Precisely what goes into this concept of ignorance? In the absence of all more concrete means of specification, there is one last fallback for rigorous and disciplined thinking, *logic*. I would argue that there are properties to ignorance, as much as there are properties to any abstraction, namely: *logical properties*. The structure and form I attribute to ignorance is not reified, because these attributes are logical. If logicality is not a sufficient condition and level of description by which to analyze an abstraction, then we are equally justified in tossing out most of the subject matter of analytic philosophy, all of pure mathematics, information science, and a good chunk of theoretical physics. Then we're justified to reject the *entire notion* of abstractions. Even if ignorance were more akin to nothingness, something can still be said of nothingness. I claim ignorance as a definite subject is amenable to concrete analysis using set-theoretic, probabilistic, and information-theoretic frameworks, and at least one form of propositional calculus, a modified or inverted version of modal epistemic logic. Ignorance also plays out in dynamical relationships with knowledge, and given that knowledge has a far more immediately salient presence, the fact of their interactivity helps ground the reality of ignorance more firmly.

It also bears mentioning that information, the theory of which has become an unstoppable juggernaut in the sciences almost immediately since its inception, is in effect a *logical* construct. Go looking for information, for information in and of itself; you will never find it. Information is a measurable logical property that is a function of an underlying system and projects off it but is not an intrinsic property of it as a substance. Ignorance—which is in fact measured in information science in the the more exact minimal definition of uncertainty, has the same kind of logical reality as information.

It is also no sign of a conspiracy that other writers (DeNicola, Rescher, Mcgoey among others) have independently derived highly similar conclusions. The literature on ignorance is still highly unconsolidated and diffuse and has not been given the same kind of institutional recognition that other fields have enjoyed. So the fact that people are coming to similar conclusions while not necessarily consciously doing so is not to be discounted as a mere coincidence.

(The author confesses, and is painfully aware of the irony, that he was unaware of the growing field of ignorance studies when he first began this project.) While there is not a perfect alignment between these analyses, it is clear to anyone who evaluates them that they approximate the same definite object. Independently derived findings, while only a heuristic, still lends support to the argument that what is being multiply noticed is real. Otherwise, one would expect each observation to vary significantly and show no correlation.

Appendix II: On Ferrier's Agnoiology

Ferrier's attempt to give a comprehensive account of metaphysics, *The Institutes of Metaphysic* included epistemology and ontology, but also included an analysis of ignorance which he rightly believed was necessary to complete the picture of philosophy. In the history of philosophy is unique in the importance it imparts to ignorance.

"The fact of our extreme ignorance being undeniable, and the science of absolute existence being apparently inaccessible except on the postulation of a universal and unlimited knowledge… Any reasoned ontological conclusion establishing what alone absolutely exists, is obviously impossible in a system which admits our ignorance without entering into any critical inquiry as to its nature; while, on the other hand, the ontology of a system which denies our ignorance, or passes it over sub silentio, must either rest upon a false ground, or upon no ground at all—on a false ground if our ignorance is denied—on no ground at all if it is not taken into account." There are traces in Ferrier's writing that hint at the substantiality of ignorance: it is not nothing. As it turns out Ferrier's theory is wrong in crucial respects. But the general theme that a theory of ignorance is needed to balance out the theory of knowledge is absolutely on point. Or as he says, "The agnoiology carries out and completes the work entered on in the epistemology."

Ferrier composed *The Institutes of Metaphysic*as a chain of connected inferences. Each antecedent proposition depends on the truth of its predecessor. Which means that if the first proposition is wrong (which it is) then the rest of

them are too. Ferrier himself admits this at the end of the section on the theory of ignorance. “In short, the overthrow of the whole agnoiology would be the consequence of the denial of the proposition which asserts that ignorance is a defect or imperfection.” Nevertheless it is instructive to review Ferrier’s assertions and why they are wrong because they still shed light on the nature of ignorance and prefigure the theory espoused in the main portion of this text.

Ferrier starts by stating that all ignorance is a defect of intelligence. It’s an imperfection of knowledge. This assertion, as is demonstrated throughout this book, is not strictly true. Ignorance—especially the act of constructively ignoring or what McGoey calls strategic ignorance—can have a use. If it can have a use in some cases, it is not a defect in all cases. Furthermore ignorance is more than the privation of knowledge, in the sense that it is not always strictly tied to a cognitive state. Ignorance is a broader concept and we can identify in the physical phenomena of nature, from the uncertainty principle, to entropy, to black holes and to the horizons of the observable universe. Because the physical universe can be thought of as an information structure, and ignorance is related to information through the concept of uncertainty, the existence of these phenomena suggest an absolute physical ignorance that is intrinsic to how nature works. We may be able to indirectly infer information about these informational lacunae based on how they influence more accessible observables, but they are at their core occluded and their inner information content is fundamentally inaccessible. What’s more, there may be absolute limits of computability as entailed by the incompleteness theorems and other theoretical considerations that place bounds on knowablity and provide an impossibility proof for the existence of a universal inference engine. It is therefore not always the case that for any piece of ignorance there is an intelligence that could potentially understand it.

This brings us to Ferrier’s second proposition. In it he claims all ignorance is remediable by *some intelligence*. To any unit of ignorance, claims Ferrier, there corresponds some intelligence, human or otherwise, capable of eliminating it. “This proposition does not prove that all ignorance is actually remedied… but only that every form of ignorance is of such a character that it may possibly be

removed; and that if certain kinds of ignorance are incident to certain orders of the intelligence, they are not, of necessity, incident to other orders of intelligence." In our terminology, this statement is only applicable to *relative ignorance:* Ignorance that relates to this or that agent's "order of intelligence", given its particular history of exposure, learnings, and capabilities. The domain of relative ignorance is vast, and might even include most unknowns. It does not include those instances of absolute ignorance about which no information can be ascertained. Nor is Ferrier right if ignorance is not strictly a defect; were it a defect, claims Ferrier, then there would in theory always exist a supplement that would correct the defect and restore it to knowledge. If ignorance is essentially a hole that means it can be filled. ("…a defect is always the privation of some quality or attribute which is consistent with the nature of the being who is deprived of it.") On my account, ignorance is instead a real thing. In information-theoretic terms, it is the objective background fuzz of uncertainty that supervenes on thermodynamic entropy. It is no more remediable than the second law of thermodynamics permits it. Because the second law places strict limits on the reversibility of physical processes, the extent to which we can piece things together is similarly restricted. Some information is lost to time because entropy has blurred it into its a equilibrated state with its surroundings. For these same reasons Ferrier's third proposition, which incidentally he calls the law of ignorance and on which he places the most significance, is also false. ("We can be ignorant only of what can possibly be known; in other words, there can be an ignorance only of that of which there can be a knowledge.")

Ferrier supposes a certain symmetry between knowledge and ignorance such that the latter can always be converted into the former. In my view ignorance is not just occluded possible knowledge. knowledge and ignorance are asymmetric, order contrasted with randomness, regularity with irregularity. knowledge is residual and arises from a process of reduction that excludes alternatives until the structure is sufficiently constrained to fit the predictable order of its object. Supposing there were this perfect parity between ignorance and knowledge, Ferrier's law of ignorance would hold. But ignorance is not merely knowledge *in potentia*. It has a completely different, opposite quality. Because ignorance is

embedded in the causal dynamics of the universe and knowledge is a byproduct of excluding the irrelevant possibilities of these dynamics, they are incommensurable.

Ferrier goes on to make an intriguing point which complicates this matter. He claims there can be no ignorance of the opposite of necessary truths. It is impossible to be ignorant of the untruths that 2+2=5 or that a space can be enclosed by two straight lines. These propositions contradict necessary truths. There is of course an infinite negative space of such anti-truths, of which we are oblivious, such as $0 > 1$, etc. Ferrier writes,

" This truth is the key to the whole philosophy of ignorance. When we consider it well, we discover that the supposition that we can be ignorant of that which is absolutely and necessarily unknowable to all intelligence, is as extreme a violation of the law of contradiction as it is possible to conceive. We perceive that a nescience of the contradictory is not ignorance, but is the very essence of intelligence; and that there can be an ignorance only of that which can be known, or, otherwise expressed, of that which is non-contradictory."

"Nescience of the contradictory," is indeed an important detail, but it begs the question as to what extent it applies. Does this mean certain limits to computability are contradictory to reason, and therefore, it's impossible to be ignorant of them because it's impossible to know them? What I consider absolute ignorance would therefore not be ignorance at all. According to Ferrier it's impossible to be ignorant of what cannot be known, but he also didn't know about the hidden regions of black holes or the quantum uncertainty principle. These constraints on knowability and the accessibility of information do not so much reflect the limitations of reason but a limit of nature which thereby imposes ultimate boundaries on reason. I cannot know both the position and velocity of a particle at the same time, not because reason won't allow it. In fact, reason tells me the opposite, I should be able to know it. Hence Richard Feynman's quip that if you think you understand quantum physics you don't understand it. Nature in its exactest details is contrary to reason. What other description of the epistemic state applies in here besides ignorance? I cannot know it, and just like the limits on provability demonstrated by the incompleteness theorems, this is consistent

with higher truths despite its contradictory relationship to reason.

Ferrier's fifth and sixth propositions also holds some interest. The fifth states, "We cannot be ignorant of any kind of objects without a subject: in other words, there can be no ignorance of objects per se, or out of relation to a mind." Ignorance indeed is a cognitive relation, and it makes little sense to talk about anything related to epistemic states without reference to a mind. To this he appends the sixth proposition, "we cannot be ignorant of material things out of all relation to a mind, subject, or self: in other words, there can be no ignorance of matter per se." This is a statement about the inscrutability of Kant's thing-in-itself. What can be known can only be known in relation to a mind, and the same can be said of ignorance. This much is true to the extent that ignorance is a purely cognitive relation. Kant's noumenon, which is precisely that object which lies outside of the categories of reason and perception, is not something of which we are ignorant, because Ferrier's definition of ignorance only admits of those quantities which are potentially knowable. In my view the noumenon is absolutely unknowable, but according to Ferrier, because the conditions of possibility for knowing it do not obtain, it is rejected altogether as an object to which there stands a cognitive relation. All conditions of knowledge and ignorance are in other words, phenomenologically bounded.

Ferrier elaborates,

"Many philosophers have seen that the human mind cannot know things by and in themselves, because it can know them only as modified and supplemented by its own faculties of cognition; in other words, that it can know them only as seen things, as touched things, and so forth—some subjective contribution being always added to the thing, and the total object apprehended being thus a composite product made up of a part which was objective, and a part which was subjective. Hence they concluded, very rashly and inconsiderately, that we were ignorant of the objective part per se, or separated from the subjective part. They adopted this counter-proposition. They gave out that we were ignorant of matter per se, of things by and in themselves. This conclusion is more particularly embraced and insisted upon by Kant."

And,

"This conclusion, however, rests on an assumption which contradicts the most strongest and essential principles of reason. It is founded on the assumption that these things may possibly be known as they are, by and in themselves, and out of relation to all intelligence… From these remarks it is obvious that Kant and other philosophers have fallen into the mistake of supposing that we could be ignorant of material things per se through an inattention to the causes which render them absolutely unknowable… They supposed that they were simply unknowable by us on account of the limitation or imperfection of our faculties of cognition, but that they were still possibly knowable by intelligences competent to know them. In the course of this work, however, it has been repeatedly shown that our incompetency to know matter per se is due to no such cause, but is attributable to the essential structure of all intelligence, and to the necessary laws of all cognition. Hence matter per se is not the simply unknowable and inconceivable to us—it is the absolutely unknowable and inconceivable in itself."

Absolute unknowability in this case according to Ferrier does not constitute absolute ignorance, which is a rather strange conclusion. It doesn't hold if we deny the proposition that only what can be possibly be known can possibly be the object of ignorance. It can't be known what lies outside of our light cone, and while in principle this information might have an intelligible form (what is outside the light cone is not unlike what is inside it, the problem is a matter of information inaccessibility), the causal dynamics strictly enforce ignorance. We can indeed be ignorant of matter, if that matter happens to lie outside our light cone. Ferrier implies that what is inconceivable to any intelligence does not constitute ignorance. Relative to some other intelligence lightyears away, the facts outside our light cone are inside theirs, and they are as structurally ignorant of us as we are them. This does not imply that there are not unintelligible facts, however. If we consider the possibility that some facts are unknowable not because of some fundamental inscrutability of their properties in relation to a mind, but because of issues of information access, then the problem becomes clearer. The applicability of his criteria are limited by the fact of meta-knowledge, we can know what we can't know. And we can know some things are impossible to know. We can realize when something is absolutely over our heads.

The primary difference between the theory of ignorance described in this text and Ferrier's is that we are able to identify an objective correlate to ignorance in information entropy. To the degree that ignorance is a subject-object cognitive relation it is bounded by physical law. The process of cognition is essentially uncertainty reduction, and it is this uncertainty which is the primary, prior substance.